



Sandia  
National  
Laboratories

[About Sandia](#)

[Capabilities](#)

[Programs](#)

[Contacting Us](#)

[News and Events](#)

[Search](#)

[Home](#)

[PDF Version  
of User's  
Guide](#)

[Zoltan Home  
Page](#)

[Zoltan  
Developer's  
Guide](#)

[How to Cite  
Zoltan](#)

# Zoltan: Data-Management Services for Parallel Applications

## User's Guide

---

### The Zoltan Team

#### Sandia National Laboratories:

[Erik Boman](#)

[Karen Devine](#)

Lee Ann Fisk

Robert Heaphy

[Bruce Hendrickson](#)

Courtenay Vaughan

### Ohio State University

[Umit Catalyurek](#)

Doruk Bozdog

### National Institute of Standards and Technology

[William F. Mitchell](#)

---

*Zoltan User's Guide, Version 2.01*

## Introduction

[Project Motivation](#)

[The Zoltan Toolkit](#)

[Terminology](#)

[Zoltan Design](#)

## Zoltan Release Notes

## Using the Zoltan Library

[System Requirements](#)

[Data Types for Object IDs](#)

[Building the Library](#)

[Testing the Library](#)

[Building Applications](#)

## **Zoltan Interface Functions**

[Error Codes](#)

[General Zoltan Interface Functions](#)

[Load-Balancing Functions](#)

[Functions for Adding Items to a Decomposition](#)

[Migration Functions](#)

[Ordering Functions](#)

[Coloring Functions](#)

## **Application-Registered Query Functions**

[General Zoltan Query Functions](#)

[Migration Query Functions](#)

## **Zoltan Parameters and Output Levels**

[General Parameters](#)

[Debugging Levels](#)

## **Load-Balancing Algorithms**

[Load-Balancing Parameters](#)

[Recursive Coordinate Bisection \(RCB\)](#)

[Recursive Inertial Bisection \(RIB\)](#)

[Hilbert Space-Filling Curve \(HSFC\) Partitioning](#)

[Refinement Tree Based Partitioning](#)

[ParMETIS](#) (graph partitioning and repartitioning)

[Jostle](#) (more graph partitioning; limited support)

[Hypergraph partitioning](#)

[Octree/Space-Filling Curve \(SFC\) Partitioning](#)

## **Ordering Algorithms**

[Nested Dissection by METIS/ParMETIS](#)

## **Coloring Algorithms**

[Parallel Coloring](#)

## **Data Services and Utilities**

[Building Utilities](#)

[Dynamic Memory Management](#)

[Unstructured Communication](#)

[Distributed Data Directories](#)

## [Examples of Library Usage](#)

[General Usage](#)

[Load-Balancing](#)

[Migration](#)

[Query Functions](#)

## [FORTRAN Interface](#)

[Compiling Zoltan](#)

[Compiling Applications](#)

[FORTRAN API](#)

[FORTRAN 77](#)

[System-Specific Remarks](#)

## [C++ Interface](#)

## [Backward Compatibility with Earlier Versions of Zoltan](#)

## [References](#)

## [Index of Interface and Query Functions](#)

---

Copyright (c) 2000,2001,2002, Sandia National Laboratories.  
The Zoltan Library and its documentation are released under the [GNU Lesser General Public License \(LGPL\)](#). See the README file in the main Zoltan directory for more information.

---

[[Zoltan Home Page](#) | [Next: Introduction](#)]

# Introduction

[Project Motivation](#)

[The Zoltan Toolkit](#)

[Terminology](#)

[Zoltan Design](#)

---

## Project Motivation

Over the past decade, parallel computers have been used with great success in many scientific simulations. While differing in their numerical methods and details of implementation, most applications successfully parallelized to date are "static" applications. Their data structures and memory usage do not change during the course of the computation. Their inter-processor communication patterns are predictable and non-varying. And their processor workloads are predictable and roughly constant throughout the simulation. Traditional finite difference and finite element methods are examples of widely used static applications.

However, increasing use of "dynamic" simulation techniques is creating new challenges for developers of parallel software. For example, adaptive finite element methods refine localized regions the mesh and/or adjust the order of the approximation on individual elements to obtain a desired accuracy in the numerical solution. As a result, memory must be allocated dynamically to allow creation of new elements or degrees of freedom. Communication patterns can vary as refinement creates new element neighbors. And localized refinement can cause severe processor load imbalance as elemental and processor work loads change throughout a simulation.

Particle simulations and crash simulations are other examples of dynamic applications. In particle simulations, scalable parallel performance depends upon a good assignment of particles to processors; grouping physically close particles within a single processor reduces inter-processor communication. Similarly, in crash simulations, assignment of physically close surfaces to a single processor enables efficient parallel contact search. In both cases, data structures and communication patterns change as particles and surfaces move. Re-partitioning of the particles or surfaces is needed to maintain geometric locality of objects within processors.

We developed the Zoltan library to simplify many of the difficulties arising in dynamic applications. Zoltan is a collection of data management services for unstructured, adaptive and dynamic applications. It includes a suite of parallel partitioning algorithms, data migration tools, distributed data directories, unstructured communication services, and dynamic memory management tools. Zoltan's data-structure neutral design allows it to be used by a variety of applications without imposing restrictions on application data structures. Its object-based interface provides a simple and inexpensive way for application developers to use the library and researchers to make new capabilities available under a common interface.

---

# The Zoltan Toolkit

The Zoltan Library contains a number of tools that simplify the development and improve the performance of parallel, unstructured and adaptive applications. The library is organized as a toolkit, so that application developers can use as little or as much of the library as desired. The major packages in Zoltan are listed below.


- A suite of [dynamic load balancing and parallel repartitioning](#) algorithms; switching between algorithms is easy, allowing straightforward comparisons of algorithms in applications.
  - [Data migration tools](#) for moving data from old partitions to new one.
  - [Distributed data directories](#): scalable (in memory and computation) algorithms for locating needed off-processor data.
  - An [unstructured communication package](#) that insulates users from the details of message sends and receives.
  - [Dynamic memory management tools](#) that greatly simplify dynamic memory debugging on state-of-the-art parallel computers.
  - A sample application [zdrive](#). It allows algorithm developers to test changes to Zoltan without having to run Zoltan in a large application code. Application developers can use the *zdrive* code to see examples of function calls to Zoltan and the implementation of query functions.
- 

## Terminology

Our design of Zoltan does not restrict it to any particular type of application. Rather, Zoltan operates on uniquely identifiable data items that we call *objects*. For example, in finite element applications, objects might be elements or nodes of the mesh. In particle applications, objects might be particles. In linear solvers, objects might be matrix rows.

Each object must have a unique *global identifier (ID)* represented as an array of unsigned integers. Common choices include global numbers of elements (nodes, particles, rows, and so on) that already exist in many applications, or a structure consisting of an owning processor number and the object's local-memory index. Objects might also have local (to a processor) IDs that do not have to be unique globally. Local IDs such as addresses or local-array indices of objects can improve the performance (and convenience) of Zoltan's interface to applications.

We use a simple example to illustrate the above terminology. In the figure [below](#), a simple finite element mesh is presented.



The blue and yellow shading indicates the mesh is partitioned for two processors. An application must provide information about the current mesh and partition to Zoltan. If, for example, the application wants Zoltan to perform operations on the elements of the mesh, it must provide information about the elements when Zoltan asks for object information.

In this example, the elements have unique numbers assigned to them, as shown by the numbers in the elements. These unique numbers can be used as global IDs in Zoltan. In addition, on each processor, local numbering information may be available. For instance, the elements owned by a processor may be stored in arrays in the processor's memory. An element's local array index may be provided to Zoltan as a local ID.

For geometric algorithms, the application must provide coordinate information to Zoltan. In this example, the coordinates of the mid-point of an element are used.

For graph-based algorithms, information about the connectivity of the objects must be provided to Zoltan. In this example, the application may consider elements connected if they share a face. The connections between elements, or *edges* of the connectivity graph, are shown in red. Connectivity information is passed to Zoltan by specifying a neighbor list for an object. The neighbor list consists of the global IDs of neighboring objects and the processor(s) currently owning those objects.

The table below summarizes the information provided to Zoltan by an application for this finite element mesh. Information about the objects includes their global and local IDs, geometry data, and graph data.

Processor	Object IDs		Geometry Data	Graph Data	
	Global	Local	(coordinates)	Neighbor Global ID List	Neighbor Processor List
Blue	1	0	(0.8,2.9)	2	Blue
	2	1	(1.7,2.9)	1,3	Blue,Blue
	3	2	(2.5,2.9)	2,4	Blue,Yellow
Yellow	4	0	(2.0,2.1)	3,5	Blue,Yellow
	5	1	(1.1,1.0)	4,6	Yellow,Yellow
	6	2	(0.5,0.2)	5,7	Yellow,Yellow

7	3	(1.3,0.2)	6,8	Yellow, Yellow
8	4	(2.1,0.2)	7	Yellow

## Zoltan Design

To make Zoltan easy to use, we do not impose any particular data structure on an application, nor do we require an application to build a particular data structure for Zoltan. Instead, Zoltan uses a [callback function interface](#), in which Zoltan queries the application for needed data. The application must provide simple functions that answer these queries.

To keep the application interface simple, we use a small set of [callback functions](#) and make them easy to write by requesting only information that is easily accessible to applications. For example, the most basic partitioning algorithms require only four callback functions. These functions return the number of objects owned by a processor, a list of weights and IDs for owned objects, the problem's dimensionality, and a given object's coordinates. More sophisticated graph-based partitioning algorithms require only two additional callback functions, which return the number of edges per object and edge lists for objects.

---

[\[Table of Contents\]](#) | [\[Next: Zoltan Release Notes\]](#) | [\[Previous: Table of Contents\]](#)

# Release Notes

Release notes are available for the following releases of Zoltan:

[Zoltan Release v2.01](#)

[Zoltan Release v2.0](#)

[Zoltan Release v1.54](#)

[Zoltan Release v1.53](#)

[Zoltan Release v1.52](#)

[Zoltan Release v1.5](#)

[Zoltan Release v1.3](#)

---

## Zoltan Release Notes v2.01

Zoltan v2.01 includes enhancements to version 2.0.

- F90 interface fixes to comply with standard F90 (e.g., shortened variable names and continuation lines). The hypergraph callback function names have changed, but C and C++ compatibility with v2.0 is maintained.
- Performance improvement to initial building of hypergraphs from application data.
- Major bug fix for dense-edge removal in parallel hypergraph method; partitioning of hypergraphs with edges containing more than 25% of the vertices was affected by this bug.
- Minor fixes to parallel hypergraph code.

---

## Zoltan Release Notes v2.0

Zoltan v2.0 includes several major additions:

- [Parallel hypergraph partitioning](#).
- [Parallel graph coloring](#), both distance-1 and distance-2.
- [Multicriteria geometric partitioning \(RCB\)](#).
- [C++ interface](#).

---

## Zoltan Release Notes v1.54

Some versions of MPICH have a bug in MPI\_Reduce\_scatter; they can report errors with MPI\_TYPE\_INDEXED. In Zoltan v1.54's unstructured communication package, calls to MPI\_Reduce\_scatter have been replaced with separate calls to MPI\_Reduce and MPI\_Scatter.



# Zoltan Release Notes v1.53

Zoltan v1.53 includes the following new capabilities:

- Portability to BSD Unix and Mac OS X was added.
  - Averaging of RCB and RIB cuts was added; see Zoltan parameter [AVERAGE\\_CUTS](#).
  - A new function [Zoltan\\_RCB\\_Box](#) returns information about subdomain bounding boxes in RCB decompositions.
  - F90 interface to [Zoltan\\_Order](#) was added.
  - Warnings that load-imbalance tolerance was not met are no longer printed when [DEBUG\\_LEVEL](#) == 0.
  - Minor bugs were addressed.
- 

# Zoltan Release Notes v1.52

Zoltan v1.52 includes the following new capabilities:

- List-based graph callback functions [ZOLTAN\\_NUM\\_EDGES\\_MULTI\\_FN](#) and [ZOLTAN\\_EDGE\\_LIST\\_MULTI\\_FN](#) were added to mirror support and performance given by the list-based geometric function [ZOLTAN\\_GEOM\\_MULTI\\_FN](#).
  - Support for ParMETIS v3.1 was added.
  - Minor bugs were addressed.
- 

# Zoltan Release Notes v1.5

This section describes improvements to Zoltan in Version 1.5. Every attempt was made to keep Zoltan v1.3 backwardly compatible with previous versions. Users of previous versions of Zoltan should refer to the [Backward Compatibility Notes](#).

Short descriptions of the following features are included below; follow the links for more details.

[Partition remapping](#)

[Unequal Numbers of Partitions and Processors](#)

[Non-Uniform Partition Sizes](#)

[Zoltan Interface Updated](#)

[Robust HSFC Box Assign](#)

[Matrix Ordering](#)

[Performance Improvements](#)

[Bug Fixes](#)

---

## Partition Remapping

During partitioning, Zoltan v1.5 can renumber partitions so that the input and output partitions have greater overlap (and, thus, lower data-migration costs). This remapping is controlled by Zoltan parameter [REMAP](#). Experiments have shown that using this parameter can greatly reduce data migration costs.

---

## Unequal Numbers of Partitions and Processors

Zoltan v1.5 can be used to generate  $k$  partitions on  $p$  processors, where  $k$  is not equal to  $p$ . Function [Zoltan\\_LB\\_Partition](#) (replacing [Zoltan\\_LB\\_Balance](#)) can generate arbitrary numbers of partitions on the given processors. The number of desired partitions is set with parameters [NUM\\_GLOBAL\\_PARTITIONS](#) or [NUM\\_LOCAL\\_PARTITIONS](#). Both partition and processor information are returned by [Zoltan\\_LB\\_Partition](#), [Zoltan\\_LB\\_Box\\_PP\\_Assign](#), and [Zoltan\\_LB\\_Point\\_PP\\_Assign](#). New Zoltan query functions [ZOLTAN\\_PARTITION\\_FN](#) and [ZOLTAN\\_PARTITION\\_MULTI\\_FN](#) return objects' partition information to Zoltan. [Zoltan\\_LB\\_Balance](#) can still be used for  $k$  equal to  $p$ .

---

## Non-Uniform Partition Sizes

Partition sizes for local and global partitions can be specified using [Zoltan\\_LB\\_Set\\_Part\\_Sizes](#), allowing non-uniformly sized partitions to be generated by Zoltan's partitioning algorithms.

---

## Zoltan Interface Updated

To support the concept of partitions separate from processors, many new interface functions were added to Zoltan v1.5 (e.g., [Zoltan\\_LB\\_Partition](#) and [Zoltan\\_Migrate](#)). These functions mimic previous Zoltan functions (e.g., [Zoltan\\_LB\\_Balance](#) and [Zoltan\\_Help\\_Migrate](#), respectively), but include both partition and processor information. Both the new and old interface functions work in Zoltan v1.5. See the notes on [Backward Compatibility](#).

---

## Robust HSFC Box Assign

Function [Zoltan\\_LB\\_Box\\_PP\\_Assign](#) now works for the [Hilbert Space-Filling Curve algorithm \(HSFC\)](#), in addition to the [RCB](#) and [RIB](#) algorithms supported in previous versions of Zoltan. [Zoltan\\_LB\\_Point\\_PP\\_Assign](#) continues to work for [HSFC](#), [RCB](#) and [RIB](#).

---

## Matrix Ordering

Zoltan v1.5 contains a matrix-ordering interface [Zoltan\\_Order](#) to ParMETIS' matrix-ordering functions. New graph-based matrix-ordering algorithms can be easily added behind this interface.

## Performance Improvements

Many performance improvements were added to Zoltan v1.5.

- List-based callback functions have been added to Zoltan ([ZOLTAN\\_GEOM\\_MULTI\\_FN](#), [ZOLTAN\\_PARTITION\\_MULTI\\_FN](#), [ZOLTAN\\_OBJ\\_SIZE\\_MULTI\\_FN](#), [ZOLTAN\\_PACK\\_OBJ\\_MULTI\\_FN](#), and [ZOLTAN\\_UNPACK\\_OBJ\\_MULTI\\_FN](#)); these functions allow entire lists of data to be passed from the application to Zoltan, replacing per-object callbacks.
  - [Zoltan\\_Migrate](#) now can accept either import lists, export lists, or both. It is no longer necessary to call [Zoltan\\_Invert\\_Lists](#) or [Zoltan\\_Compute\\_Destinations](#) to get appropriate input for [Zoltan\\_Migrate](#).
  - Zoltan v1.5 contains performance improvements within individual algorithms. We recommend users upgrade to the latest version.
- 

## Bug Fixes

Bug fixes were made to Zoltan's algorithms and interface. Users of previous versions of Zoltan are encouraged to upgrade.

---

# Zoltan Release Notes v1.3

This section describes improvements to Zoltan in Version 1.3. Every attempt was made to keep Zoltan v1.3 backwardly compatible with previous versions. Users of previous versions of Zoltan should refer to the [Backward Compatibility Notes](#).

Short descriptions of the following features are included below; follow the links for more details.

[More Data Services](#)

[New Hilbert Space-Filling Curve Partitioning](#)

[Support for Structured-Grid Partitioning](#)

[Support for ParMETIS v3.0](#)

[Performance Improvements](#)

[Zoltan Interface Updated](#)

[Improved Test Suite](#)

[Bug Fixes](#)

---

## More Data Services

Zoltan's mission has been widened beyond its original focus on dynamic load-balancing algorithms. Now Zoltan also provides data management services to parallel, unstructured, and adaptive computations. Several packages of parallel data services have been added and made available to application developers.

These services include the following:

- An [unstructured communication package](#) that simplifies complicated communication by insulating applications from the details of message sends and receives.
  - A [distributed data directory](#) that allows applications to efficiently (in memory and time) locate off-processor data.
  - A [dynamic memory management package](#) that simplifies debugging of memory allocation problems on state-of-the-art parallel computers.
- 

## New Hilbert Space-Filling Curve Partitioning

Zoltan now includes a fast, efficient implementation of [Hilbert Space-Filling Curve \(HSFC\)](#) partitioning. This geometric method also includes support for [Zoltan\\_LB\\_Box\\_Assign](#) and [Zoltan\\_LB\\_Point\\_Assign](#) functions.

---

## Support for Structured-Grid Partitioning

Zoltan's [Recursive Coordinate Bisection \(RCB\)](#) partitioning algorithm has been enhanced to allow generation of strictly rectilinear subdomains. This capability can be used for partitioning of grids for structured-grid applications. See parameter [RCB\\_RECTILINEAR\\_BLOCKS](#).

---

## Support for ParMETIS v3.0

In addition to providing interfaces to [ParMETIS v2.0](#) and [PJostle](#), Zoltan now provides an interfaces [ParMETIS v3.0](#). Full support of ParMETIS v3.0's multiconstraint and multiobjective partitioning is included.

---

## Performance Improvements

Performance of Zoltan's partitioning algorithms has been improved through a number of code optimizations and new features. In addition, user parameter [RETURN\\_LISTS](#) can be used to specify which returned arguments are computed by [Zoltan\\_LB\\_Balance](#), allowing reduced work in partitioning. In the [Recursive Coordinate Bisection \(RCB\)](#) partitioning algorithm, user parameters allow cut directions to be locked in an attempt to minimize data movement; see parameters [RCB\\_LOCK DIRECTIONS](#) and [RCB\\_SET DIRECTIONS](#).

---

## Zoltan Interface Updated

Zoltan has adopted a more modular design, making it easier to use by applications and easier to modify by algorithm developers. Names in the [Zoltan interface](#) and code are tied more closely to their functionality. Full [backward compatibility](#) is supported for users of previous versions of Zoltan.

---

## Improved Test Suite

The Zoltan [test suite](#) has been improved, with more tests providing greater code coverage and platform-specific answer files accounting for differences due to computer architectures.

---

## Bug Fixes

Some bug fixes were made to Zoltan's algorithms and interface. Users of previous versions of Zoltan are encouraged to upgrade.

---

[\[Table of Contents](#) | [Next: Using the Zoltan Library](#) | [Previous: Introduction\]](#)

# Using the Zoltan library

This section contains information needed to use the Zoltan library with applications:

[System requirements.](#)

[Data types](#) for global and local IDs.

Instructions for [building the Zoltan library.](#)

Instructions for [building applications](#) that use Zoltan.

## System Requirements

Zoltan was designed to run on parallel computers and clusters of workstations. In order to build and use Zoltan, you will need:

- ANSI C compiler.
- [MPI](#) library for message passing (version 1.1 or higher), such as MPICH or LAM.
- A Unix-like operating system (e.g., Linux or Solaris) and *gmake* (GNU Make) are recommended to build the library.
- A Fortran90 compatible compiler is required if you wish to [use Zoltan with Fortran applications.](#)

Zoltan has been tested on a variety of platforms, including Linux, Solaris, Irix, and the [ASCI Red](#) Teraflop machine. If you wish to use Zoltan on a non-Unix operating system, for example Windows NT or 2000, you will have to port Zoltan yourself.

## Data Types for Object IDs

Application query functions and application callable library functions use global and local identifiers (IDs) for objects. *All objects to be used in load balancing must have unique global IDs.* Zoltan stores an ID as an array of unsigned integers. The number of entries in these arrays can be set using the [NUM\\_GID\\_ENTRIES](#) and [NUM\\_LID\\_ENTRIES](#) parameters; by default, one unsigned integer represents an ID. Applications may use whatever format is most convenient to store their IDs; the IDs can then be converted to and from Zoltan's ID format in the [application-registered query functions](#).

The following type definitions are defined in *include/zoltan\_types.h*; they can be used by an application for memory allocation, MPI communication, and as arguments to [load-balancing interface functions](#) and [application-registered query functions](#).

```
typedef unsigned int ZOLTAN_ID_TYPE;  
typedef ZOLTAN_ID_TYPE *ZOLTAN_ID_PTR;  
#define ZOLTAN_ID_MPI_TYPE MPI_UNSIGNED
```

In the Fortran interface, IDs are passed as arrays of integers since unsigned integers are not supported in Fortran. See the description of the [Fortran interface](#) for more details.

The local IDs passed to Zoltan are not used by the library; they are provided for the convenience of the

application and can contain any information desired by the application. For instance, local array indices for objects may be passed as local IDs, enabling direct access to object data in the query function routines. See the [application-registered query functions](#) for more details. The source code distribution contains an example application [zdrive](#) in which global IDs are integers and local IDs are local array indices. One may choose not to use local ids at all, in which case [NUM\\_LID\\_ENTRIES](#) may be set to zero.

Some Zoltan routines (e.g., [Zoltan\\_LB\\_Partition](#) and [Zoltan\\_Invert\\_Lists](#)) allocate arrays of type **ZOLTAN\_ID\_PTR** and return them to the application. Others (e.g., **Zoltan\_Order** and [Zoltan\\_DD\\_Find](#)) require the application to allocate memory for IDs. Memory for IDs can be allocated as follows:

```
ZOLTAN_ID_PTR gids;
int num_gids, int num_gid_entries;
gids = (ZOLTAN_ID_PTR) ZOLTAN\_MALLOC(num_gids *
num_gid_entries * sizeof(ZOLTAN_ID_TYPE);
```

The system call *malloc* may be used instead of [ZOLTAN\\_MALLOC](#).

## Building the Zoltan Library

The Zoltan library is implemented in ANSI C and can be compiled with any ANSI C compiler. Makefiles are included with the source code; these makefiles require the GNU Make (gmake) utility. The top-level Makefile defines targets for the Zoltan library, [test driver](#) programs in C, C++ and Fortran90, and [two graphical utilities](#) useful for visualization of geometric partitions. (The test drivers and utilities are primarily intended for use by developers.) This Makefile need not be edited to build Zoltan. Instead, environment-specific definitions are specified in the configuration file, *Utilities/Config/Config.<platform>*, where *<platform>* specifies the particular platform for which Zoltan is being built. Paths to compilers, include files, and libraries are defined in this file and are then read by the top-level Makefile. Examples of configuration files for Solaris, Sandia's ASCI Red (tflop) computer, SGI workstations, and PCs running Linux are included in the *Utilities/Config* subdirectory. A well-commented version of the configuration file, *Utilities/Config/Config.generic*, is also included; this file can be used as a template for new environment-specific files. The variables in these files should be edited to reflect the new system's environment.

The command for building Zoltan is shown below:

```
gmake [options] zoltan
```

where the options that may be specified are listed below.

### Options to gmake:

**ZOLTAN\_ARCH=<platform>** Specify the target architecture for the Zoltan library. A corresponding file, *Utilities/Config/Config.<platform>*, containing environment definitions for *<platform>*, must be created in the *Utilities/Config* directory.



**`YES_FORTRAN=1`**

Include Fortran support in the Zoltan library. By default, the Zoltan library is built without the interface that allows use from Fortran applications. If this option is specified, the [Fortran interface](#) is compiled and included in the library. Use of this option requires that a Fortran 90 (or 95, or later) compiler is available.

As an alternative to typing the options on the gmake command line, they may be set as environment variables; e.g., if you are using a C-shell (csh or tcsh), type

```
setenv ZOLTAN_ARCH <platform>
```

or if you are using a Bourne-type shell (e.g., sh or bash), type

```
ZOLTAN_ARCH = <platform>; export ZOLTAN_ARCH
```

The resulting library *libzoltan.a*, object files, and dependency files are stored in the directory `Obj_<platform>`.

## Testing the Zoltan Library

The *examples* directory contains simple C and C++ examples which use the Zoltan library. The Makefile in this directory has three targets:

```
gmake ZOLTAN_ARCH=<platform> C_Examples
```

This builds simple C language examples that use the Zoltan library to perform load balancing.

```
gmake ZOLTAN_ARCH=<platform> CPP_Examples
```

This builds simple C++ language examples that use the Zoltan library to perform load balancing. To build C++ applications, define CPPC to point to your C++ compiler in the `Config.<platform>` file.

```
gmake ZOLTAN_ARCH=<platform> all
```

Build both C and C++ examples. Don't forget to define CPPC in your `Config.<platform>` file.

Some of these examples make use of a small library of support routines found in the *examples/lib* directory. These routines create simple test meshes of varying sizes, perform error checking across the parallel application, and define Zoltan call backs.

The "right" answer for these tests depends on the number of processes with which you run the tests. In general, if they compile successfully, run quickly (in seconds), and produce reasonable looking output, then you have been successful in building Zoltan.

## Building Applications that use Zoltan

The C library interface is described in the include file *include/zoltan.h*; this file should be included in all C application source files that call Zoltan library routines.

The [C++ interface](#) to Zoltan is implemented in header files which define classes that wrap the Zoltan C library. The file *include/zoltan\_cpp.h* defines the **Zoltan** class which encapsulates a load balancing data structure and the Zoltan load balancing functions which operate upon it. Include this header file instead in your C++ application. Note that C++ applications should call the C function **Zoltan\_Initialize** before



creating a **Zoltan** object.

[Fortran applications](#) must USE [module zoltan](#) and specify *Zoltan/Obj\_<platform>* as a directory to be searched for module information files.

The C, C++ or Fortran application should then be linked with the Zoltan library (built with Fortran support in the Fortran case) and its [utility libraries](#) by including

*-lzoltan*

in the linking command for the application. Communication within Zoltan is performed through MPI, so appropriate MPI libraries must be linked with the application. Third-party libraries, such as [ParMETIS](#) and [Jostle](#), must be also be linked with the application if they were included in compilation of the Zoltan library. (A courtesy copy of [ParMETIS](#) is included with the Zoltan distribution; [Jostle](#) must be obtained directly from <http://www.gre.ac.uk/~jjg01/>.)

For applications that used versions of Zoltan before Zoltan v.1.3, only minor updates to the application build process are needed; see the section on [backward compatibility](#) of Zoltan.

---

[\[Table of Contents\]](#) | [Next: Zoltan Interface Functions](#) | [Previous: Zoltan Release Notes](#)

# Zoltan Interface Functions

An application calls a series of dynamic load-balancing library functions to initialize the load balancer, perform load balancing and migrate data. This section describes the syntax of each type of interface function:

[General Zoltan Interface Functions](#)  
[Load-Balancing Interface Functions](#)  
[Functions for Augmenting a Decomposition](#)  
[Migration Interface Functions](#)

Examples of the calling sequences for initialization, load-balancing, and data migration are included in the [Initialization](#), [Load-Balancing](#), and [Migration](#) sections, respectively, of the [Examples of Library Usage](#).

---

## Error Codes

All interface functions, with the exception of [Zoltan\\_Create](#), return an error code to the application. The possible return codes are defined in *include/zoltan\_types.h* and Fortran [module zoltan](#), and are listed in the [table](#) below.

Note: Robust error handling in parallel has not yet been achieved in Zoltan. When a processor returns from Zoltan due to an error condition, other processors do not necessarily return the same condition. In fact, other processors may not know that the original processor has returned from Zoltan, and may wait indefinitely in a communication routine (e.g., waiting for a message from the original processor that is not sent due to the error condition). The parallel error-handling capabilities of Zoltan will be improved in future releases.

<i>ZOLTAN_OK</i>	Function returned without warnings or errors.
<i>ZOLTAN_WARN</i>	Function returned with warnings. The application will probably be able to continue to run.
<i>ZOLTAN_FATAL</i>	A fatal error occurred within the Zoltan library.
<i>ZOLTAN_MEMERR</i>	An error occurred while allocating memory. When this error occurs, the library frees any allocated memory and returns control to the application. If the application then wants to try to use another, less memory-intensive algorithm, it can do so.

*Return codes defined in include/zoltan\_types.h.*

---

## Naming conventions

The C, Fortran and C++ interfaces follow consistent naming conventions, as illustrated in the following table.

	C and Fortran	C++
Partitioning and migration functions example: perform partitioning example: assign a point to a partition	<b>Zoltan_function()</b> <b>Zoltan_LB_Partition()</b> <b>Zoltan_LB_Point_Assign()</b>	<b>Zoltan::function()</b> <b>Zoltan::LB_Partition()</b> <b>Zoltan::LB_Point_Assign()</b>
Unstructured communication example: perform communication	<b>Zoltan_Comm_function</b> <b>Zoltan_Comm_Do()</b>	<b>Zoltan_Comm::function</b> <b>Zoltan_Comm::Do()</b>
Distributed data example: find objects in a remote process	<b>Zoltan_DD_function</b> <b>Zoltan_DD_Find()</b>	<b>Zoltan_DD::function</b> <b>Zoltan_DD::Find()</b>
Timers example: print timing results	<b>Zoltan_Timer_function</b> <b>Zoltan_Timer_Print()</b>	<b>Zoltan_Timer::function</b> <b>Zoltan_Timer::Print()</b>

In particular, the C++ **Zoltan** class represents a load balancing instance and the methods that operate on it. The method name is identical to the part of the C and Fortran function name that indicates the function performed. A C++ **Zoltan\_Comm** object represents an instance of unstructured communication, a C++ **Zoltan\_DD** object represents a distributed directory, and a C++ **Zoltan\_Timer** object is a timer. Their method names are derived similarly.

---

[[Table of Contents](#) | [Next: Initialization Functions](#) | [Previous: Using the Library](#)]

# General Interface Functions

Functions used to initialize and manipulate Zoltan's data structures are described below:

[\*\*Zoltan\\_Initialize\*\*](#)

[\*\*Zoltan\\_Create\*\*](#)

[\*\*Zoltan\\_Copy\*\*](#)

[\*\*Zoltan\\_Copy\\_To\*\*](#)

[\*\*Zoltan\\_Set\\_Param\*\*](#)

[\*\*Zoltan\\_Set\\_Param\\_Vec\*\*](#)

[\*\*Zoltan\\_Set\\_Fn\*\*](#)

[\*\*Zoltan\\_Set\\_<zoltan\\_fn\\_type>\\_Fn\*\*](#)

[\*\*Zoltan\\_Destroy\*\*](#)

---

C and C++:     int **Zoltan\_Initialize** (  
                  int *argc*,  
                  char \*\**argv*,  
                  float \**ver*);

FORTTRAN:     FUNCTION **Zoltan\_Initialize**( *argc*, *argv*, *ver*)  
                  INTEGER(Zoltan\_INT) :: Zoltan\_Initialize  
                  INTEGER(Zoltan\_INT), INTENT(IN), OPTIONAL :: *argc*  
                  CHARACTER(LEN=\*), DIMENSION(\*), INTENT(IN), OPTIONAL :: *argv*  
                  REAL(Zoltan\_FLOAT), INTENT(OUT) :: *ver*

---

The **Zoltan\_Initialize** function initializes MPI for Zoltan. If the application uses MPI, this function should be called after calling **MPI\_Init**. If the application does not use MPI, this function calls **MPI\_Init** for use by Zoltan. This function is called with the *argc* and *argv* command-line arguments from the main program, which are used if **Zoltan\_Initialize** calls **MPI\_Init**. From C, if **MPI\_Init** has already been called, the *argc* and *argv* arguments may have any value because their values will be ignored. From Fortran, if one of *argc* or *argv* is omitted, they must both be omitted. If they are omitted, *ver* does NOT have to be passed as a keyword argument.

**Zoltan\_Initialize** returns the Zoltan version number so that users can verify which version of the library their application is linked to.

C++ applications should call the C **Zoltan\_Initialize** function before using the C++ interface to the Zoltan library.

## Arguments:

<i>argc</i>	The number of command-line arguments to the application.
<i>argv</i>	An array of strings containing the command-line arguments to the application.
<i>ver</i>	Upon return, the version number of the library.

**Returned Value:**

int                    [Error code](#).

```

C:          struct Zoltan_Struct *Zoltan_Create (
            MPI_Comm communicator);
FORTRAN:    FUNCTION Zoltan_Create(communicator)
            TYPE(Zoltan_Struct), pointer :: Zoltan_Create
            INTEGER, INTENT(IN) :: communicator
C++:        Zoltan (
            const MPI_Comm &communicator = MPI_COMM_WORLD);

```

The **Zoltan\_Create** function allocates memory for storage of information to be used by Zoltan and sets the default values for the information. The pointer returned by this function is passed to many subsequent functions. An application may allocate more than one **Zoltan\_Struct** data structure; for example, an application may use several **Zoltan\_Struct** structures if, say, it uses different decompositions with different load-balancing techniques.

In the C++ interface to Zoltan, the **Zoltan** class represents a Zoltan load balancing data structure and the functions that operate on it. It is the constructor which allocates an instance of a **Zoltan** object. It has no return value.

**Arguments:**

*communicator*      The MPI communicator to be used for this Zoltan structure. Only those processors included in the communicator participate in Zoltan functions. If all processors are to participate, *communicator* should be **MPI\_COMM\_WORLD** .

**Returned Value:**

struct **Zoltan\_Struct** \*      Pointer to memory for storage of Zoltan information. If an error occurs, NULL will be returned in C, or the result will be a nullified pointer in Fortran. Any error that occurs in this function is assumed to be fatal.

```

C:          struct Zoltan_Struct *Zoltan_Copy (
            Zoltan_Struct *from);
FORTRAN:    FUNCTION Zoltan_Copy(from)
            TYPE(Zoltan_Struct), pointer :: Zoltan_Copy
            TYPE(Zoltan_Struct), INTENT(IN) :: from
C++:        Zoltan (
            const Zoltan &zz);

```

The **Zoltan\_Copy** function creates a new **Zoltan\_Struct** and copies the state of the existing **Zoltan\_Struct**, which it has been passed, to the new structure. It returns the new **Zoltan\_Struct**.

There is no direct interface to **Zoltan\_Copy** from C++. Rather, the **Zoltan** copy constructor invokes the C library **Zoltan\_Copy** program.

**Arguments:**

*from*                      A pointer to the **Zoltan\_Struct** that is to be copied.

**Returned Value:**

struct  
**Zoltan\_Struct** \*        Pointer to a new **Zoltan\_Struct**, which is now a copy of *from*.

```
C:          int Zoltan_Copy_To (
              Zoltan_Struct *to,
              Zoltan_Struct *from);

FORTRAN:    FUNCTION Zoltan_Copy_To(to, from)
              INTEGER(Zoltan_INT) :: Zoltan_Copy_To
              TYPE(Zoltan_Struct), INTENT(IN) :: to
              TYPE(Zoltan_Struct), INTENT(IN) :: from

C++:        Zoltan & operator= (
              const Zoltan &zz);
```

The **Zoltan\_Copy\_To** function copies one **Zoltan\_Struct** to another, after first freeing any memory used by the target **Zoltan\_Struct** and re-initializing it.

The C++ interface to the **Zoltan\_Copy\_To** function is through the **Zoltan** copy operator, which invokes the C library **Zoltan\_Copy\_To** program.

**Arguments:**

*to*                        A pointer to an existing **Zoltan\_Struct**, the target of the copy.  
*from*                      A pointer to an existing **Zoltan\_Struct**, the source of the copy.

**Returned Value:**

int                        **0** on success and **1** on failure.

```
C:          int Zoltan_Set_Param (
              struct Zoltan_Struct *zz,
              char *param_name,
              char *new_val);

FORTRAN:    FUNCTION Zoltan_Set_Param(zz, param_name, new_val)
              INTEGER(Zoltan_INT) :: Zoltan_Set_Param
              TYPE(Zoltan_Struct), INTENT(IN) :: zz
              CHARACTER(LEN=*), INTENT(IN) :: param_name, new_value

C++:        int Zoltan::Set_Param (
              const std::string &param_name,
              const std::string &new_value);
```

**Zoltan\_Set\_Param** is used to alter the value of one of the parameters used by Zoltan. All Zoltan

parameters have reasonable default values, but this routine allows a user to provide alternative values if desired.

### Arguments:

<i>zz</i>	Pointer to the Zoltan structure created by <a href="#">Zoltan_Create</a> .
<i>param_name</i>	A string containing the name of the parameter to be altered. Note that the string is case-insensitive. Also, different Zoltan structures can have different parameter values.
<i>new_val</i>	A string containing the new value for the parameter. Example strings include "3.154", "True", "7" or anything appropriate for the parameter being set. As above, the string is case-insensitive.

### Returned Value:

int [Error code](#).

---



---

```

C:      int Zoltan_Set_Param_Vec (
        struct Zoltan_Struct *zz,
        char *param_name,
        char *new_val,
        int index);

FORTRAN:  FUNCTION Zoltan_Set_Param_Vec(zz, param_name, new_val, index)
          INTEGER(Zoltan_INT) :: Zoltan_Set_Param_Vec
          TYPE(Zoltan_Struct), INTENT(IN) :: zz
          CHARACTER(LEN=*), INTENT(IN) :: param_name, new_value
          INTEGER(Zoltan_INT), INTENT(IN) :: index

C++:      int Zoltan::Set_Param_Vec (
        const std::string &param_name,
        const std::string &new_val,
        const int &index);
  
```

---

**Zoltan\_Set\_Param\_Vec** is used to alter the value of a vector parameter in Zoltan. A vector parameter is a parameter that has one name but contains multiple values. These values are referenced by their indices, usually starting at 0. Each entry (component) may have a different value. This routine sets a single entry (component) of a vector parameter. If you want all entries (components) of a vector parameter to have the same value, set the parameter using [Zoltan\\_Set\\_Param](#) as if it were a scalar parameter. If one only sets the values of a subset of the indices for a vector parameter, the remaining entries will have the default value for that particular parameter.

### Arguments:

<i>zz</i>	Pointer to the Zoltan structure created by <a href="#">Zoltan_Create</a> .
<i>param_name</i>	A string containing the name of the parameter to be altered. Note that the string is case-insensitive. Also, different Zoltan structures can have different parameter values.

<i>new_val</i>	A string containing the new value for the parameter. Example strings include "3.154", "True", "7" or anything appropriate for the parameter being set. As above, the string is case-insensitive.
<i>index</i>	The index of the entry of the vector parameter to be set. The default in Zoltan is that the first entry in a vector has index 0 (C-style indexing).

**Returned Value:**

int      [Error code](#).

C:            int **Zoltan\_Set\_Fn** (  
               struct **Zoltan\_Struct** \*zz,  
               **ZOLTAN\_FN\_TYPE** fn\_type,  
               void (\*fn\_ptr)(),  
               void \*data);

FORTRAN:    FUNCTION **Zoltan\_Set\_Fn**(zz, fn\_type, fn\_ptr, data)  
               INTEGER(Zoltan\_INT) :: Zoltan\_Set\_Fn  
               TYPE(Zoltan\_Struct), INTENT(IN) :: zz  
               TYPE(ZOLTAN\_FN\_TYPE), INTENT(IN) :: fn\_type  
               EXTERNAL :: fn\_ptr  
               <type-data>, OPTIONAL :: data

<type-data> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE), DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where x is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

C++:        int **Zoltan::Set\_Fn** (  
               const **ZOLTAN\_FN\_TYPE** &fn\_type,  
               void (\*fn\_ptr)(),  
               void \*data = 0);

**Zoltan\_Set\_Fn** registers an application-supplied query function in the Zoltan structure. All types of query functions can be registered through calls to **Zoltan\_Set\_Fn**. To register functions while maintaining strict type-checking of the *fn\_ptr* argument, use [Zoltan\\_Set\\_<zoltan\\_fn\\_type>\\_Fn](#).

**Arguments:**

<i>zz</i>	Pointer to the Zoltan structure created by <a href="#">Zoltan_Create</a> .
<i>fn_type</i>	The type of function being registered; see <a href="#">Application-Registered Query Functions</a> for possible function types.
<i>fn_ptr</i>	A pointer to the application-supplied query function being registered.
<i>data</i>	A pointer to user defined data that will be passed, as an argument, to the function pointed to by <i>fn_ptr</i> . In C it may be NULL. In Fortran it may be omitted.

**Returned Value:**

int      [Error code](#).



C:           int **Zoltan\_Set\_**<[zoltan\\_fn\\_type](#)>\_Fn (  
               struct **Zoltan\_Struct** \*zz,  
               <[zoltan\\_fn\\_type](#)> (\*fn\_ptr)(),  
               void \*data);

FORTRAN:   FUNCTION **Zoltan\_Set\_**<[zoltan\\_fn\\_type](#)>\_Fn(zz, fn\_ptr, data)  
               INTEGER(Zoltan\_INT) :: Zoltan\_Set\_<[zoltan\\_fn\\_type](#)>\_Fn  
               TYPE(Zoltan\_Struct), INTENT(IN) :: zz  
               EXTERNAL :: fn\_ptr  
               <type-data>, OPTIONAL :: data

An interface block for *fn\_ptr* is included in the FUNCTION definition so that strict type-checking of the registered query function can be done.

<type-data> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE), DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

C++:       int **Zoltan::Set\_**<[zoltan\\_fn\\_type](#)>\_Fn (  
               <[zoltan\\_fn\\_type](#)> (\*fn\_ptr)(),  
               void \*data = 0);

---

The interface functions **Zoltan\_Set\_**<[zoltan\\_fn\\_type](#)>\_Fn, where <[zoltan\\_fn\\_type](#)> is one of the query function types, register specific types of [application-supplied query functions](#) in the Zoltan structure. One interface function exists for each type of query function. For example, **Zoltan\_Set\_Num\_Geom\_Fn** registers a query function of type [ZOLTAN\\_NUM\\_GEOM\\_FN](#). Each query function has an associated **Zoltan\_Set\_**<[zoltan\\_fn\\_type](#)>\_Fn. A complete list of these functions is included in *include/zoltan.h*.

Query functions can be registered using either [Zoltan\\_Set\\_Fn](#) or **Zoltan\_Set\_**<[zoltan\\_fn\\_type](#)>\_Fn. **Zoltan\_Set\_**<[zoltan\\_fn\\_type](#)>\_Fn provides strict type checking of the *fn\_ptr* argument; the argument's type is specified for each **Zoltan\_Set\_**<[zoltan\\_fn\\_type](#)>\_Fn. [Zoltan\\_Set\\_Fn](#) does not provide this strict type checking, as the pointer to the registered function is cast to a void pointer.

### Arguments:

<i>zz</i>	Pointer to the Zoltan structure created by <a href="#">Zoltan_Create</a> .
<i>fn_ptr</i>	A pointer to the application-supplied query function being registered. The type of the pointer matches < <a href="#">zoltan_fn_type</a> > in the name <b>Zoltan_Set_</b> < <a href="#">zoltan_fn_type</a> >_Fn.
<i>data</i>	A pointer to user defined data that will be passed, as an argument, to the function pointed to by <i>fn_ptr</i> . In C it may be NULL. In Fortran it may be omitted.

### Returned Value:

int	<a href="#">Error code</a> .
-----	------------------------------

### Example:

The interface function

```
int Zoltan_Set_Geom_Fn(struct Zoltan_Struct *zz, ZOLTAN\_GEOM\_FN
(*fn_ptr)(),
                        void *data);
```

registers an [ZOLTAN\\_GEOM\\_FN](#) query function.

---

```
C:      void Zoltan_Destroy (
        struct Zoltan_Struct **zz);
FORTRAN: SUBROUTINE Zoltan_Destroy(zz)
        TYPE(Zoltan_Struct), POINTER :: zz
C++:    ~Zoltan ();
```

---

**Zoltan\_Destroy** frees the memory associated with a Zoltan structure and sets the structure to NULL in C or nullifies the structure in Fortran. Note that **Zoltan\_Destroy** does not deallocate the import and export arrays returned from Zoltan (e.g., the arrays returned from [Zoltan\\_LB\\_Partition](#)); these arrays can be deallocated through a separate call to [Zoltan\\_LB\\_Free\\_Part](#).

There is no explicit **Destroy** method in the C++ interface. The **Zoltan** object is destroyed when the destructor executes.

As a side effect, **Zoltan\_Destroy** (and the C++ **Zoltan** destructor) frees the MPI communicator that had been allocated for the structure. So it is important that the application does not call **MPI\_Finalize** before it calls **Zoltan\_Destroy** or before the destructor executes.

### Arguments:

<i>zz</i>	A pointer to the address of the Zoltan structure, created by <a href="#">Zoltan_Create</a> , to be destroyed.
-----------	---

---

[\[Table of Contents\]](#) | [\[Next: Load-Balancing Functions\]](#) | [\[Previous: Zoltan Interface Functions\]](#)

# Load-Balancing Functions

The following functions are the load-balancing interface functions in the Zoltan library; their descriptions are included below.

[Zoltan\\_LB\\_Partition](#)

[Zoltan\\_LB\\_Set\\_Part\\_Sizes](#)

[Zoltan\\_LB\\_Eval](#)

[Zoltan\\_LB\\_Free\\_Part](#)

For [backward compatibility](#) with previous versions of Zoltan, the following functions are also maintained. These functions are applicable only when the number of partitions to be generated is equal to the number of processors on which the partitions are computed. That is, these functions assume "partitions" and "processors" are synonymous.

[Zoltan\\_LB\\_Balance](#)

[Zoltan\\_LB\\_Free\\_Data](#)

Descriptions of algorithm-specific interface functions are included with the documentation of their associated algorithms. Algorithm-specific functions include:

[Zoltan\\_RCB\\_Box](#)

---

---

```

C:      int Zoltan_LB_Partition (
          struct Zoltan_Struct *zz,
          int *changes,
          int *num_gid_entries,
          int *num_lid_entries,
          int *num_import,
          ZOLTAN\_ID\_PTR *import_global_ids,
          ZOLTAN\_ID\_PTR *import_local_ids,
          int **import_procs,
          int **import_to_part,
          int *num_export,
          ZOLTAN\_ID\_PTR *export_global_ids,
          ZOLTAN\_ID\_PTR *export_local_ids,
          int **export_procs,
          int **export_to_part);

```

```

FORTRAN:  FUNCTION Zoltan_LB_Partition(zz, changes, num_gid_entries, num_lid_entries,
      num_import, import_global_ids, import_local_ids, import_procs, import_to_part,
      num_export, export_global_ids, export_local_ids, export_procs, export_to_part)
      INTEGER(Zoltan_INT) :: Zoltan_LB_Partition
      TYPE(Zoltan_Struct), INTENT(IN) :: zz
      LOGICAL, INTENT(OUT) :: changes
      INTEGER(Zoltan_INT), INTENT(OUT) :: num_gid_entries, num_lid_entries
      INTEGER(Zoltan_INT), INTENT(OUT) :: num_import, num_export
      INTEGER(Zoltan_INT), POINTER, DIMENSION(:) :: import_global_ids,
      export_global_ids
      INTEGER(Zoltan_INT), POINTER, DIMENSION(:) :: import_local_ids,
      export_local_ids
      INTEGER(Zoltan_INT), POINTER, DIMENSION(:) :: import_procs, export_procs
      INTEGER(Zoltan_INT), POINTER, DIMENSION(:) :: import_to_part, export_to_part

C++:      int Zoltan::LB_Partition (
      int &changes,
      int &num_gid_entries,
      int &num_lid_entries,
      int &num_import,
      ZOLTAN\_ID\_PTR &import_global_ids,
      ZOLTAN\_ID\_PTR &import_local_ids,
      int * &import_procs,
      int * &import_to_part,
      int &num_export,
      ZOLTAN\_ID\_PTR &export_global_ids,
      ZOLTAN\_ID\_PTR &export_local_ids,
      int * &export_procs,
      int * &export_to_part);

```

**Zoltan\_LB\_Partition** invokes the load-balancing routine specified by the [LB\\_METHOD](#) parameter. The number of partitions it generates is specified by the [NUM\\_GLOBAL\\_PARTITIONS](#) or [NUM\\_LOCAL\\_PARTITIONS](#) parameters. Results of the partitioning are returned in lists of objects to be imported into and exported from partitions on this processor. Objects are included in these lists if *either* their partition assignment or their processor assignment is changed by the new decomposition. If an application requests multiple partitions on a single processor, these lists may include objects whose partition assignment is changing, but whose processor assignment is unchanged.

Returned arrays are allocated in Zoltan; applications should not allocate these arrays before calling **Zoltan\_LB\_Partition**. The arrays are later freed through calls to [Zoltan\\_LB\\_Free\\_Part](#).

### Arguments:

*zz*                      Pointer to the Zoltan structure, created by [Zoltan\\_Create](#), to be used in this invocation of the load-balancing routine.

<i>changes</i>	Set to 1 or .TRUE. if the decomposition was changed by the load-balancing method; 0 or .FALSE. otherwise.
<i>num_gid_entries</i>	Upon return, the number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	Upon return, the number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>num_import</i>	Upon return, the number of objects that are newly assigned to this processor or to partitions on this processor (i.e., the number of objects being imported from different partitions to partitions on this processor). If the value returned is -1, no import information has been returned and all import arrays below are NULL. (The <a href="#">RETURN_LISTS</a> parameter determines whether import lists are returned).
<i>import_global_ids</i>	Upon return, an array of <i>num_import</i> global IDs of objects to be imported to partitions on this processor. (size = <i>num_import</i> * <i>num_gid_entries</i> )
<i>import_local_ids</i>	Upon return, an array of <i>num_import</i> local IDs of objects to be imported to partitions on this processor. (size = <i>num_import</i> * <i>num_lid_entries</i> )
<i>import_procs</i>	Upon return, an array of size <i>num_import</i> listing the processor IDs of the processors that owned the imported objects in the previous decomposition (i.e., the source processors).
<i>import_to_part</i>	Upon return, an array of size <i>num_import</i> listing the partitions to which the imported objects are being imported.
<i>num_export</i>	Upon return, this value of this count and the following lists depends on the value of the <a href="#">RETURN_LISTS</a> parameter: <ul style="list-style-type: none"> <li>● It is the count of objects on this processor that are newly assigned to other processors or to other partitions on this processor, if <a href="#">RETURN_LISTS</a> is "EXPORT" or "EXPORT AND IMPORT".</li> <li>● It is the count of all objects on this processor, if <a href="#">RETURN_LISTS</a> is "PARTITION ASSIGNMENTS".</li> <li>● It is -1 if the value of <a href="#">RETURN_LISTS</a> indicates that either no lists are to be returned, or only import lists are to be returned. If the value returned is -1, no export information has been returned and all export arrays below are NULL .</li> </ul>
<i>export_global_ids</i>	Upon return, an array of <i>num_export</i> global IDs of objects to be exported from partitions on this processor (if <a href="#">RETURN_LISTS</a> is equal to "EXPORT" or "EXPORT AND IMPORT"), or an array of <i>num_export</i> global IDs for every object on this processor (if <a href="#">RETURN_LISTS</a> is equal to "PARTITION ASSIGNMENTS"), . (size = <i>num_export</i> * <i>num_gid_entries</i> )

<i>export_local_ids</i>	Upon return, an array of <i>num_export</i> local IDs associated with the global IDs returned in <i>export_global_ids</i> (size = <i>num_export</i> * <i>num_lid_entries</i> )
<i>export_procs</i>	Upon return, an array of size <i>num_export</i> listing the processor ID of the processor to which each object is now assigned (i.e., the destination processor). If <a href="#">RETURN_LISTS</a> is equal to "PARTITION ASSIGNMENTS", this list includes all objects, otherwise it only includes the objects which are moving to a new partition and/or process.
<i>export_to_part</i>	Upon return, an array of size <i>num_export</i> listing the partitions to which the objects are assigned under the new partitioning.

**Returned Value:**

int            [Error code](#).

```

C:          int Zoltan_LB_Set_Part_Sizes (
            struct Zoltan_Struct *zz,
            int global_num,
            int len,
            int *part_ids,
            int *wgt_idx,
            float *part_sizes);

FORTRAN:    function Zoltan_LB_Set_Part_Sizes( zz,global_part,len,partids,wgtidx,partsizes)
            integer(Zoltan_INT) :: Zoltan_LB_Set_Part_Sizes
            type(Zoltan_Struct) INTENT(IN) zz
            integer(Zoltan_INT) INTENT(IN) global_part,len,partids(*),wgtidx(*)
            real(Zoltan_FLOAT) INTENT(IN) partsizes(*)

C++:        int Zoltan::LB_Set_Part_Sizes (
            const int &global_num,
            const int &len,
            int *part_ids,
            int *wgt_idx,
            float *part_sizes);

```

**Zoltan\_LB\_Set\_Part\_Sizes** is used to specify the desired partition sizes in Zoltan. By default, Zoltan assumes that all partitions should be of equal size. With **Zoltan\_LB\_Set\_Part\_Sizes**, one can specify the relative (not absolute) sizes of the partitions. For example, if two partitions are requested and the desired sizes are 1 and 2, that means that the first partition will be assigned approximately one third of the total load. If the sizes were instead given as 1/3 and 2/3, respectively, the result would be exactly the same. Note that if there are multiple weights per object, one can (must) specify the partition size for each weight dimension independently.

**Arguments:**

zz            Pointer to the Zoltan structure created by [Zoltan\\_Create](#).

<i>global_num</i>	Set to 1 if global partition numbers are given, 0 otherwise (local partition numbers).
<i>len</i>	Length of the next three input arrays.
<i>part_ids</i>	Array of partition numbers, either global or local. (Partition numbers are integers starting from 0.)
<i>vwgt_idx</i>	Array of weight indices (between 0 and OBJ_WEIGHT_DIM-1). This array should contain all zeros when there is only one weight per object.
<i>part_sizes</i>	Relative values for partition sizes; <i>part_sizes[i]</i> is the desired relative size of the <i>vwgt_idx[i]</i> 'th weight of partition <i>part_ids[i]</i> .

**Returned Value:**

int                    [Error code](#).

```

C:      int Zoltan_LB_Eval (
        struct Zoltan_Struct *zz,
        int print_stats,
        int *nobj,
        float *obj_wgt,
        int *ncuts,
        float *cut_wgt,
        int *nboundary,
        int *nadj);

FORTRAN: FUNCTION Zoltan_LB_Eval(zz, print_stats, nobj, obj_wgt, ncuts, cut_wgt,
nboundary, nadj)
INTEGER(Zoltan_INT) :: Zoltan_LB_Eval
TYPE(Zoltan_Struct), INTENT(IN) :: zz
LOGICAL, INTENT(IN) :: print_stats
INTEGER(Zoltan_INT), INTENT(OUT), OPTIONAL :: nobj, ncuts, nboundary, nadj
REAL(Zoltan_FLOAT), INTENT(OUT), DIMENSION(vwgt_dim), OPTIONAL ::
obj_wgt
REAL(Zoltan_FLOAT), INTENT(OUT), DIMENSION(ewgt_dim), OPTIONAL ::
cut_wgt

C++:    int Zoltan::LB_Eval (
        const int &print_stats,
        int *nobj,
        float * const obj_wgt,
        int *ncuts,
        float * const cut_wgt,
        int *nboundary,
        int *nadj);

```

**Zoltan\_LB\_Eval** evaluates the quality of a decomposition. Some quality metrics are available only if the graph query functions have been registered. **Zoltan\_LB\_Eval** may either print a summary of the results to *stdout* or return the results in the output parameters. *NOTE:* The interface to this function may change in future versions of Zoltan. Users are discouraged from relying on the output arguments from



## Zoltan\_LB\_Eval.

**Arguments:**

<i>zz</i>	Pointer to the Zoltan structure.
<i>print_stats</i>	If <i>print_stats</i> >0 (.TRUE. in Fortran), print a summary (max, min, and sum) of the quality metrics to <i>stdout</i> .
<i>nobj</i>	Upon return, the number of objects on this processor.
<i>obj_wgt</i>	Upon return, an array (of dimension <a href="#">OBJ_WEIGHT_DIM</a> ) containing the sum of object weights on this processor.
<i>ncuts</i>	Upon return, the number of (communication) edge cuts for this processor.
<i>cut_wgt</i>	Upon return, an array (of dimension <a href="#">EDGE_WEIGHT_DIM</a> ) of cut weights for this processor.
<i>nboundary</i>	Upon return, the number of boundary objects on this processor.
<i>nadj</i>	Upon return, the number of adjacent processors as defined by the communication graph.

**Returned Value:**

int      [Error code.](#)

**Query functions:**

Required:      [ZOLTAN\\_NUM\\_OBJ\\_FN](#)  
                   [ZOLTAN\\_OBJ\\_LIST\\_FN](#) or  
                   [ZOLTAN\\_FIRST\\_OBJ\\_FN/ZOLTAN\\_NEXT\\_OBJ\\_FN](#)

Optional:      [ZOLTAN\\_NUM\\_EDGES\\_MULTI\\_FN](#) or [ZOLTAN\\_NUM\\_EDGES\\_FN](#)  
                   [ZOLTAN\\_EDGE\\_LIST\\_MULTI\\_FN](#) or [ZOLTAN\\_EDGE\\_LIST\\_FN](#)

An output parameter is returned only if the input value of that parameter was not NULL. The rationale for this feature is that if one wishes just to print the evaluation results, one can simply set all (or some of) the output parameters to NULL in the function call. From Fortran, one may omit one or more of the optional output parameters.

Note that the sum of *ncuts* over all processors is actually twice the number of edges cut in the graph (because each edge is counted twice). The same principle holds for *cut\_wgt*.

There are a few improvements in Zoltan\_LB\_Eval in Zoltan version 1.5 (or higher). First, the balance data are computed with respect to both processors and partitions (if applicable). Second, the desired partition sizes (as set by Zoltan\_LB\_Set\_Partition\_Sizes) are taken into account when computing the imbalance.

Known bug: If a partition is spread across several processors, the computed cut information (*ncuts* and *cut\_wgt*) may be incorrect (too high).



```

C:      int Zoltan_LB_Free_Part (
        ZOLTAN\_ID\_PTR *global_ids,
        ZOLTAN\_ID\_PTR *local_ids,
        int **procs,
        int **to_part);

FORTRAN: FUNCTION Zoltan_LB_Free_Part(global_ids, local_ids, procs, to_part)
INTEGER(Zoltan_INT) :: Zoltan_LB_Free_Part
INTEGER(Zoltan_INT), POINTER, DIMENSION(:) :: global_ids
INTEGER(Zoltan_INT), POINTER, DIMENSION(:) :: local_ids
INTEGER(Zoltan_INT), POINTER, DIMENSION(:) :: procs, to_part

C++:    int Zoltan::LB_Free_Part (
        ZOLTAN\_ID\_PTR *global_ids,
        ZOLTAN\_ID\_PTR *local_ids,
        int **procs,
        int **to_part);

```

---

**Zoltan\_LB\_Free\_Part** frees the memory allocated by Zoltan to return the results of [Zoltan\\_LB\\_Partition](#) or [Zoltan\\_Invert\\_Lists](#). Memory pointed to by the arguments is freed and the arguments are set to NULL in C and C++ or nullified in Fortran. NULL arguments may be passed to **Zoltan\_LB\_Free\_Part**. Note that this function does not destroy the Zoltan data structure itself; it is deallocated through a call to [Zoltan\\_Destroy](#) in C and Fortran and by the object destructor in C++.

### Arguments:

<i>global_ids</i>	An array containing the global IDs of objects.
<i>local_ids</i>	An array containing the local IDs of objects.
<i>procs</i>	An array containing processor IDs.
<i>to_part</i>	An array containing partition numbers.

### Returned Value:

int	<a href="#">Error code</a> .
-----	------------------------------

---

```

C:      int Zoltan_LB_Balance (
        struct Zoltan_Struct *zz,
        int *changes,
        int *num_gid_entries,
        int *num_lid_entries,
        int *num_import,
        ZOLTAN\_ID\_PTR *import_global_ids,
        ZOLTAN\_ID\_PTR *import_local_ids,
        int **import_procs,
        int *num_export,
        ZOLTAN\_ID\_PTR *export_global_ids,
        ZOLTAN\_ID\_PTR *export_local_ids,

```

```
int **export_procs);
```

**FORTTRAN:**    **FUNCTION Zoltan\_LB\_Balance**(*zz, changes, num\_gid\_entries, num\_lid\_entries, num\_import, import\_global\_ids, import\_local\_ids, import\_procs, num\_export, export\_global\_ids, export\_local\_ids, export\_procs*)  
 INTEGER(Zoltan\_INT) :: Zoltan\_LB\_Balance  
 TYPE(Zoltan\_Struct), INTENT(IN) :: *zz*  
 LOGICAL, INTENT(OUT) :: *changes*  
 INTEGER(Zoltan\_INT), INTENT(OUT) :: *num\_gid\_entries, num\_lid\_entries*  
 INTEGER(Zoltan\_INT), INTENT(OUT) :: *num\_import, num\_export*  
 INTEGER(Zoltan\_INT), POINTER, DIMENSION(:) :: *import\_global\_ids, export\_global\_ids*  
 INTEGER(Zoltan\_INT), POINTER, DIMENSION(:) :: *import\_local\_ids, export\_local\_ids*  
 INTEGER(Zoltan\_INT), POINTER, DIMENSION(:) :: *import\_procs, export\_procs*

---

**Zoltan\_LB\_Balance** is a wrapper around [Zoltan\\_LB\\_Partition](#) that excludes the partition assignment results. **Zoltan\_LB\_Balance** assumes the number of partitions is equal to the number of processors; thus, the partition assignment is equivalent to the processor assignment. Results of the partitioning are returned in lists of objects to be imported and exported. These arrays are allocated in Zoltan; applications should not allocate these arrays before calling **Zoltan\_LB\_Balance**. The arrays are later freed through calls to [Zoltan\\_LB\\_Free\\_Data](#) or [Zoltan\\_LB\\_Free\\_Part](#).

### Arguments:

All arguments are analogous to those in [Zoltan\\_LB\\_Partition](#).

Partition-assignment arguments *import\_to\_part* and *export\_to\_part* are not included, as processor and partitions numbers are considered to be the same in **Zoltan\_LB\_Balance**.

### Returned Value:

int            [Error code](#).

---

**C:**            int **Zoltan\_LB\_Free\_Data** (  
               [ZOLTAN\\_ID\\_PTR](#) \**import\_global\_ids*,  
               [ZOLTAN\\_ID\\_PTR](#) \**import\_local\_ids*,  
               int \*\**import\_procs*,  
               [ZOLTAN\\_ID\\_PTR](#) \**export\_global\_ids*,  
               [ZOLTAN\\_ID\\_PTR](#) \**export\_local\_ids*,  
               int \*\**export\_procs*);

```

FORTRAN:  FUNCTION Zoltan_LB_Free_Data(import_global_ids, import_local_ids,
      import_procs, export_global_ids, export_local_ids, export_procs)
      INTEGER(Zoltan_INT) :: Zoltan_LB_Free_Data
      INTEGER(Zoltan_INT), POINTER, DIMENSION(:) :: import_global_ids,
      export_global_ids
      INTEGER(Zoltan_INT), POINTER, DIMENSION(:) :: import_local_ids,
      export_local_ids
      INTEGER(Zoltan_INT), POINTER, DIMENSION(:) :: import_procs, export_procs

```

---

**Zoltan\_LB\_Free\_Data** frees the memory allocated by the Zoltan to return the results of [Zoltan\\_LB\\_Balance](#) or [Zoltan\\_Compute\\_Destinations](#). Memory pointed to by the arguments is freed and the arguments are set to NULL in C or nullified in Fortran. NULL arguments may be passed to **Zoltan\_LB\_Free\_Data**. Note that this function does not destroy the Zoltan data structure itself; it is deallocated through a call to [Zoltan\\_Destroy](#).

### Arguments:

<i>import_global_ids</i>	The array containing the global IDs of objects imported to this processor.
<i>import_local_ids</i>	The array containing the local IDs of objects imported to this processor.
<i>import_procs</i>	The array containing the processor IDs of the processors that owned the imported objects in the previous decomposition (i.e., the source processors).
<i>export_global_ids</i>	The array containing the global IDs of objects exported from this processor.
<i>export_local_ids</i>	The array containing the local IDs of objects exported from this processor.
<i>export_procs</i>	The array containing the processor IDs of processors that own the exported objects in the new decomposition (i.e., the destination processors).

### Returned Value:

int      [Error code](#).

---

[\[Table of Contents\]](#) | [Next: Functions for Augmenting a Decomposition](#) | [Previous: Initialization Functions](#)

# Functions for Augmenting a Decomposition

The following functions support the addition of new items to an existing decomposition. Given a decomposition, they determine to which processor(s) a new item should be assigned. Currently, they work in conjunction with only the [RCB](#), [RIB](#), and [HSFC](#) algorithms.

[Zoltan\\_LB\\_Point\\_PP\\_Assign](#)

[Zoltan\\_LB\\_Box\\_PP\\_Assign](#)

For [backward compatibility](#) with previous versions of Zoltan, the following functions are also maintained. These functions are applicable only when the number of partitions to be generated is equal to the number of processors on which the partitions are computed. That is, these functions assume "partitions" and "processors" are synonymous.

[Zoltan\\_LB\\_Point\\_Assign](#)

[Zoltan\\_LB\\_Box\\_Assign](#)

---

```

C:          int Zoltan_LB_Point_PP_Assign (
              struct Zoltan_Struct * zz,
              double * coords,
              int * proc,
              int * part );

FORTRAN:    FUNCTION Zoltan_LB_Point_PP_Assign(zz, coords, proc, part)
              INTEGER(Zoltan_INT) :: Zoltan_LB_Point_PP_Assign
              TYPE(Zoltan_Struct), INTENT(IN) :: zz
              REAL(Zoltan_DOUBLE), DIMENSION(*), INTENT(IN) :: coords
              INTEGER(Zoltan_INT), INTENT(OUT) :: proc
              INTEGER(Zoltan_INT), INTENT(OUT) :: part

C++:        int Zoltan::LB_Point_PP_Assign (
              double * const coords,
              int & proc,
              int & part );
  
```

---

**Zoltan\_LB\_Point\_PP\_Assign** is used to determine to which processor and partition a new point should be assigned. It is applicable only to geometrically generated decompositions ([RCB](#), [RIB](#), and [HSFC](#)). If the parameter **KEEP\_CUTS** is set to TRUE, then the sequence of cuts that define the decomposition is saved. Given a new geometric point, the processor and partition which own it can be determined.

## Arguments:

<i>zz</i>	Pointer to the Zoltan structure created by <a href="#">Zoltan_Create</a> .
<i>coords</i>	The (x,y) or (x,y,z) coordinates of the point being assigned.
<i>proc</i>	Upon return, the ID of the processor to which the point should belong.
<i>part</i>	Upon return, the ID of the partition to which the point should belong.

**Returned Value:**

int                    [Error code](#).

---

```

C:      int Zoltan_LB_Box_PP_Assign (
        struct Zoltan_Struct * zz,
        double xmin,
        double ymin,
        double zmin,
        double xmax,
        double ymax,
        double zmax,
        int *procs,
        int *numprocs,
        int *parts,
        int *numparts);

FORTRAN: FUNCTION Zoltan_LB_Box_PP_Assign(zz, xmin, ymin, zmin, xmax, ymax, zmax,
      procs, numprocs, parts, numparts)
      INTEGER(Zoltan_INT) :: Zoltan_LB_Box_PP_Assign
      TYPE(Zoltan_Struct), INTENT(IN) :: zz
      REAL(Zoltan_DOUBLE), INTENT(IN) :: xmin, ymin, zmin, xmax, ymax, zmax
      INTEGER(Zoltan_INT), DIMENSION(*), INTENT(OUT) :: procs
      INTEGER(Zoltan_INT), INTENT(OUT) :: numprocs
      INTEGER(Zoltan_INT), DIMENSION(*), INTENT(OUT) :: parts
      INTEGER(Zoltan_INT), INTENT(OUT) :: numparts

C++:    int Zoltan::LB_Box_PP_Assign (
        const double & xmin,
        const double & ymin,
        const double & zmin,
        const double & xmax,
        const double & ymax,
        const double & zmax,
        int * const procs,
        int & numprocs,
        int * const parts,
        int & numparts);

```

---

In many settings, it is useful to know which processors and partitions might need to know about an extended geometric object. **Zoltan\_LB\_Box\_PP\_Assign** addresses this problem. Given a geometric decomposition of space (currently only [RCB](#), [RIB](#), and [HSFC](#) are supported), and given an axis-aligned box around the geometric object, **Zoltan\_LB\_Box\_PP\_Assign** determines which processors and partitions own geometry that intersects the box. To use this routine, the parameter **KEEP\_CUTS** must be set to TRUE when the decomposition is generated. This parameter will cause the sequence of geometric cuts to be saved, which is necessary for **Zoltan\_LB\_Box\_PP\_Assign** to do its job.

Note that if the parameter **REDUCE\_DIMENSIONS** was set to TRUE and the geometry was

determined to be degenerate when decomposition was calculated, then the calculation was performed on transformed coordinates. This means that **Zoltan\_LB\_Box\_PP\_Assign** must transform the supplied bounding box accordingly. The transformed vertices are bounded again, and the partition intersections are calculated in the transformed space on this new bounding box. The impact of this is that **Zoltan\_LB\_Box\_PP\_Assign** may return partitions not actually intersecting the original bounding box, but it will not omit any partitions intersecting the original bounding box.

### Arguments:

<i>zz</i>	Pointer to the Zoltan structure created by <a href="#">Zoltan_Create</a> .
<i>xmin, ymin, zmin</i>	The coordinates of the lower extent of the bounding box around the object. If the geometry is two-dimensional, the <i>z</i> value is ignored.
<i>xmax, ymax, zmax</i>	The coordinates of the upper extent of the bounding box around the object. If the geometry is two-dimensional, the <i>z</i> value is ignored.
<i>procs</i>	The list of processors intersecting the box are returned starting at this address. Note that <i>it is the responsibility of the calling routine to ensure that there is sufficient space for the return list</i> .
<i>numprocs</i>	Upon return, this value contains the number of processors that intersect the box (i.e. the number of entries placed in the <i>procs</i> list).
<i>parts</i>	The list of partitions intersecting the box are returned starting at this address. Note that <i>it is the responsibility of the calling routine to ensure that there is sufficient space for the return list</i> .
<i>numparts</i>	Upon return, this value contains the number of partitions that intersect the box (i.e. the number of entries placed in the <i>parts</i> list).

### Returned Value:

int [Error code](#).

---

```

C:      int Zoltan_LB_Point_Assign (
        struct Zoltan_Struct * zz,
        double * coords,
        int * proc);

FORTRAN: FUNCTION Zoltan_LB_Point_Assign(zz, coords, proc)
        INTEGER(Zoltan_INT) :: Zoltan_LB_Point_Assign
        TYPE(Zoltan_Struct), INTENT(IN) :: zz
        REAL(Zoltan_DOUBLE), DIMENSION(*), INTENT(IN) :: coords
        INTEGER(Zoltan_INT), INTENT(OUT) :: proc

```

---

**Zoltan\_LB\_Point\_Assign** is a wrapper around [Zoltan\\_LB\\_Point\\_PP\\_Assign](#) that excludes the partition assignment results. **Zoltan\_LB\_Point\_Assign** assumes the number of partitions is equal to the number of processors; thus, the partition assignment is equivalent to the processor assignment.

### Arguments:

All arguments are analogous to those in [Zoltan\\_LB\\_Point\\_PP\\_Assign](#).

Partition-assignment argument *part* is not included, as processor and partitions numbers are considered to be the same in **Zoltan\_LB\_Point\_Assign**.

### Returned Value:

int      [Error code](#).

C:            int **Zoltan\_LB\_Box\_Assign** (  
               struct **Zoltan\_Struct** \* zz,  
               double *xmin*,  
               double *ymin*,  
               double *zmin*,  
               double *xmax*,  
               double *ymax*,  
               double *zmax*,  
               int \**procs*,  
               int \**numprocs*);

FORTTRAN:   FUNCTION **Zoltan\_LB\_Box\_Assign**(zz, *xmin*, *ymin*, *zmin*, *xmax*, *ymax*, *zmax*, *procs*,  
                  *numprocs*)  
               INTEGER(Zoltan\_INT) :: Zoltan\_LB\_Box\_Assign  
               TYPE(Zoltan\_Struct), INTENT(IN) :: zz  
               REAL(Zoltan\_DOUBLE), INTENT(IN) :: *xmin*, *ymin*, *zmin*, *xmax*, *ymax*, *zmax*  
               INTEGER(Zoltan\_INT), DIMENSION(\*), INTENT(OUT) :: *procs*  
               INTEGER(Zoltan\_INT), INTENT(OUT) :: *numprocs*

**Zoltan\_LB\_Box\_Assign** is a wrapper around [Zoltan\\_LB\\_Box\\_PP\\_Assign](#) that excludes the partition assignment results. **Zoltan\_LB\_Box\_Assign** assumes the number of partitions is equal to the number of processors; thus, the partition assignment is equivalent to the processor assignment.

### Arguments:

All arguments are analogous to those in [Zoltan\\_LB\\_Box\\_PP\\_Assign](#).

Partition-assignment arguments *parts* and *numparts* are not included, as processor and partitions numbers are considered to be the same in **Zoltan\_LB\_Box\_Assign**.

### Returned Value:

int      [Error code](#).

[[Table of Contents](#) | [Next: Migration Functions](#) | [Previous: Load-Balancing Functions](#)]



# Migration Functions

Zoltan's migration functions transfer object data to the processors in a new decomposition. Data to be transferred is specified through the import/export lists returned by [Zoltan\\_LB\\_Partition](#) (or [Zoltan\\_LB\\_Balance](#)). Alternatively, users may specify their own import/export lists.

The migration functions can migrate objects based on their new partition assignments and/or their new processor assignments. Behavior is determined by the [MIGRATE\\_ONLY\\_PROC\\_CHANGES](#) parameter.

If requested, Zoltan can automatically transfer an application's data between processors to realize a new decomposition. This functionality will be performed as part of the call to [Zoltan\\_LB\\_Partition](#) (or [Zoltan\\_LB\\_Balance](#)) if the [AUTO\\_MIGRATE](#) parameter is set to TRUE (nonzero) via a call to [Zoltan\\_Set\\_Param](#). This approach is effective for when the data to be moved is relatively simple. For more complicated data movement, the application can leave [AUTO\\_MIGRATE](#) FALSE and call [Zoltan\\_Migrate](#) (or [Zoltan\\_Help\\_Migrate](#)) itself. In either case, [routines to pack and unpack object data](#) must be provided by the application. See the [Migration Examples](#) for examples with and without auto-migration.

The following functions are the migration interface functions. Their detailed descriptions can be found below.

[Zoltan\\_Invert\\_Lists](#)

[Zoltan\\_Migrate](#)

The following functions are maintained for [backward compatibility](#) with previous versions of Zoltan. These functions are applicable only when the number of partitions to be generated is equal to the number of processors on which the partitions are computed. That is, these functions assume "partitions" and "processors" are synonymous.

[Zoltan\\_Compute\\_Destinations](#)

[Zoltan\\_Help\\_Migrate](#)

---

---



```

C:      int Zoltan_Invert_Lists (
        struct Zoltan_Struct *zz,
        int num_known,
        ZOLTAN\_ID\_PTR known_global_ids,
        ZOLTAN\_ID\_PTR known_local_ids,
        int *known_procs,
        int *known_to_part,
        int *num_found,
        ZOLTAN\_ID\_PTR *found_global_ids,
        ZOLTAN\_ID\_PTR *found_local_ids,
        int **found_procs,
        int **found_to_part);

FORTRAN: FUNCTION Zoltan_Invert_Lists(zz, num_known, known_global_ids,
    known_local_ids, known_procs, known_to_part, num_found, found_global_ids,
    found_local_ids, found_procs, found_to_part)
    INTEGER(Zoltan_INT) :: Zoltan_Invert_Lists
    TYPE(Zoltan_Struct),INTENT(IN) :: zz
    INTEGER(Zoltan_INT), INTENT(IN) :: num_known
    INTEGER(Zoltan_INT), INTENT(OUT) :: num_found
    INTEGER(Zoltan_INT), POINTER, DIMENSION(:) :: known_global_ids,
    found_global_ids
    INTEGER(Zoltan_INT), POINTER, DIMENSION(:) :: known_local_ids,
    found_local_ids
    INTEGER(Zoltan_INT), POINTER, DIMENSION(:) :: known_procs, found_procs
    INTEGER(Zoltan_INT), POINTER, DIMENSION(:) :: known_to_part, found_to_part

C++:    int Zoltan::Invert_Lists (
        const int & num_known,
        ZOLTAN\_ID\_PTR const known_global_ids,
        ZOLTAN\_ID\_PTR const known_local_ids,
        int * const known_procs,
        int * const known_to_part,
        int &num_found,
        ZOLTAN\_ID\_PTR &found_global_ids,
        ZOLTAN\_ID\_PTR &found_local_ids,
        int * &found_procs,
        int * &found_to_part);

```

---

**Zoltan\_Invert\_Lists** computes inverse communication maps useful for migrating data. It can be used in two ways:

- Given a list of known off-processor objects to be received by a processor, compute a list of local objects to be sent by the processor to other processors; or
- Given a list of known local objects to be sent by a processor to other processors, compute a list of off-processor objects to be received by the processor.

For example, if each processor knows which objects it will import from other processors,

**Zoltan\_Invert\_Lists** computes the list of objects each processor needs to export to other processors. If, instead, each processor knows which objects it will export to other processors, **Zoltan\_Invert\_Lists** computes the list of objects each processor will import from other processors. The computed lists are allocated in Zoltan; they should not be allocated by the application before calling **Zoltan\_Invert\_Lists**. These lists can be freed through a call to [Zoltan\\_LB\\_Free\\_Part](#).

### Arguments:

<i>zz</i>	Pointer to the Zoltan structure, created by <a href="#">Zoltan_Create</a> , to be used in this invocation of the migration routine.
<i>num_known</i>	The number of known objects to be received (sent) by this processor.
<i>known_global_ids</i>	An array of <i>num_known</i> global IDs of known objects to be received (sent) by this processor. (size = <i>num_known</i> * <a href="#">NUM_GID_ENTRIES</a> )
<i>known_local_ids</i>	An array of <i>num_known</i> local IDs of known objects to be received (sent) by this processor. (size = <i>num_known</i> * <a href="#">NUM_LID_ENTRIES</a> )
<i>known_procs</i>	An array of size <i>num_known</i> listing the processor IDs of the processors that the known objects will be received from (sent to).
<i>known_to_part</i>	An array of size <i>num_known</i> listing the partition numbers of the partitions that the known objects will be assigned to.
<i>num_found</i>	Upon return, the number of objects that must be sent to (received from) other processors.
<i>found_global_ids</i>	Upon return, an array of <i>num_found</i> global IDs of objects to be sent (received) by this processor. (size = <i>num_found</i> * <a href="#">NUM_GID_ENTRIES</a> )
<i>found_local_ids</i>	Upon return, an array of <i>num_found</i> local IDs of objects to be sent (received) by this processor. (size = <i>num_found</i> * <a href="#">NUM_LID_ENTRIES</a> )
<i>found_procs</i>	Upon return, an array of size <i>num_found</i> listing the processor IDs of processors that the found objects will be sent to (received from).
<i>found_to_part</i>	An array of size <i>num_found</i> listing the partition numbers of the partitions that the found objects will be assigned to.

### Returned Value:

int      [Error code](#).

Note that the number of global and local ID entries ([NUM\\_GID\\_ENTRIES](#) and [NUM\\_LID\\_ENTRIES](#)) should be set using [Zoltan\\_Set\\_Param](#) before calling **Zoltan\_Invert\_Lists**. All processors must have the same values for these two parameters.

```

C:      int Zoltan_Migrate (
        struct Zoltan_Struct *zz,
        int num_import,
        ZOLTAN\_ID\_PTR import_global_ids,
        ZOLTAN\_ID\_PTR import_local_ids,
        int *import_procs,
        int *import_to_part,
        int num_export,
        ZOLTAN\_ID\_PTR export_global_ids,
        ZOLTAN\_ID\_PTR export_local_ids,
        int *export_procs,
        int *export_to_part);

FORTRAN: FUNCTION Zoltan_Migrate(zz, num_import, import_global_ids, import_local_ids,
import_procs, import_to_part, num_export, export_global_ids, export_local_ids,
export_procs, export_to_part)
INTEGER(Zoltan_INT) :: Zoltan_Migrate
TYPE(Zoltan_Struct),INTENT(IN) :: zz
INTEGER(Zoltan_INT), INTENT(IN) :: num_import, num_export
INTEGER(Zoltan_INT), POINTER, DIMENSION(:) :: import_global_ids,
export_global_ids
INTEGER(Zoltan_INT), POINTER, DIMENSION(:) :: import_local_ids,
export_local_ids
INTEGER(Zoltan_INT), POINTER, DIMENSION(:) :: import_procs, export_procs
INTEGER(Zoltan_INT), POINTER, DIMENSION(:) :: import_to_part, export_to_part

C++:    int Zoltan::Migrate (
        const int & num_import,
        ZOLTAN\_ID\_PTR const import_global_ids,
        ZOLTAN\_ID\_PTR const import_local_ids,
        int * const import_procs,
        int * const import_to_part,
        const int & num_export,
        ZOLTAN\_ID\_PTR const export_global_ids,
        ZOLTAN\_ID\_PTR const export_local_ids,
        int * const export_procs,
        int * const export_to_part);

```

**Zoltan\_Migrate** takes lists of objects to be sent to other processors, along with the destinations of those objects, and performs the operations necessary to send the data associated with those objects to their destinations. **Zoltan\_Migrate** performs the following operations using the application-registered functions:

- Call [ZOLTAN\\_PRE\\_MIGRATE\\_PP\\_FN\\_TYPE](#) (if registered)
- For each export object, call [ZOLTAN\\_OBJ\\_SIZE\\_FN\\_TYPE](#) to get object sizes.
- For each export object, call [ZOLTAN\\_PACK\\_OBJ\\_FN\\_TYPE](#) to load communication buffers.

- Communicate buffers to destination processors.
- Call [ZOLTAN\\_MID\\_MIGRATE\\_PP\\_FN\\_TYPE](#) (if registered).
- For each imported object, call [ZOLTAN\\_UNPACK\\_OBJ\\_FN\\_TYPE](#) to move data from the buffer into the new processor's data structures.
- Call [ZOLTAN\\_POST\\_MIGRATE\\_PP\\_FN\\_TYPE](#) (if registered).

Either export lists or import lists must be specified for **Zoltan\_Migrate**. Both export lists and import lists may be specified, but both are not required.

If export lists are provided, non-NULL values for input arguments *import\_global\_ids*, *import\_local\_ids*, *import\_procs*, and *import\_to\_part* are optional. The values must be non-NULL only if no export lists are provided or if the import lists are used by the application callback functions [ZOLTAN\\_PRE\\_MIGRATE\\_PP\\_FN](#), [ZOLTAN\\_MID\\_MIGRATE\\_PP\\_FN](#), and [ZOLTAN\\_POST\\_MIGRATE\\_PP\\_FN](#). If all processors pass NULL arguments for the import arrays, the value of *num\_import* should be -1.

Similarly, if import lists are provided, non-NULL values for input arguments *export\_global\_ids*, *export\_local\_ids*, *export\_procs*, and *export\_to\_part* are optional. The values must be non-NULL only if no import lists are provided or if the export lists are used by the application callback functions [ZOLTAN\\_PRE\\_MIGRATE\\_PP\\_FN](#), [ZOLTAN\\_MID\\_MIGRATE\\_PP\\_FN](#), and [ZOLTAN\\_POST\\_MIGRATE\\_PP\\_FN](#). If all processors pass NULL arguments for the export arrays, the value of *num\_export* should be -1. In this case, **Zoltan\_Migrate** computes the export lists based on the import lists.

### Arguments:

<i>zz</i>	Pointer to the Zoltan structure, created by <a href="#"><u>Zoltan_Create</u></a> , to be used in this invocation of the migration routine.
<i>num_import</i>	The number of objects to be imported to partitions on this processor; these objects may be stored on other processors or may be moving to new partitions within this processor. Use <i>num_import</i> =-1 if all processors do not specify import arrays.
<i>import_global_ids</i>	An array of <i>num_import</i> global IDs of objects to be imported to partitions on this processor. (size = <i>num_import</i> * <a href="#"><u>NUM_GID_ENTRIES</u></a> ). All processors may pass <i>import_global_ids</i> =NULL if export lists are provided and <i>import_global_ids</i> is not needed by callback functions <a href="#"><u>ZOLTAN_PRE_MIGRATE_PP_FN</u></a> , <a href="#"><u>ZOLTAN_MID_MIGRATE_PP_FN</u></a> , and <a href="#"><u>ZOLTAN_POST_MIGRATE_PP_FN</u></a> .

- import\_local\_ids* An array of *num\_import* local IDs of objects to be imported to partitions on this processor.  
(size = *num\_import* \* [NUM\\_LID\\_ENTRIES](#))  
All processors may pass *import\_local\_ids*=NULL if export lists are provided and *import\_local\_ids* is not needed by callback functions  
[ZOLTAN\\_PRE\\_MIGRATE\\_PP\\_FN](#), [ZOLTAN\\_MID\\_MIGRATE\\_PP\\_FN](#),  
and [ZOLTAN\\_POST\\_MIGRATE\\_PP\\_FN](#).
- import\_procs* An array of size *num\_import* listing the processor IDs of objects to be imported to partitions on this processor (i.e., the source processors).  
All processors may pass *import\_procs*=NULL if export lists are provided and *import\_procs* is not needed by callback functions  
[ZOLTAN\\_PRE\\_MIGRATE\\_PP\\_FN](#), [ZOLTAN\\_MID\\_MIGRATE\\_PP\\_FN](#),  
and [ZOLTAN\\_POST\\_MIGRATE\\_PP\\_FN](#).
- import\_to\_part* An array of size *num\_import* listing the partitions to which imported objects should be assigned.  
All processors may pass *import\_to\_part*=NULL if export lists are provided and *import\_to\_part* is not needed by callback functions  
[ZOLTAN\\_PRE\\_MIGRATE\\_PP\\_FN](#), [ZOLTAN\\_MID\\_MIGRATE\\_PP\\_FN](#),  
and [ZOLTAN\\_POST\\_MIGRATE\\_PP\\_FN](#).
- num\_export* The number of objects that were stored on this processor in the previous decomposition that are assigned to other processors or to different partitions within this processor in the new decomposition.  
Use *num\_export*=-1 if all processors do not specify export arrays.
- export\_global\_ids* An array of *num\_export* global IDs of objects to be exported to new partitions.  
(size = *num\_export* \* [NUM\\_GID\\_ENTRIES](#))  
All processors may pass *export\_global\_ids*=NULL if import lists are provided and *export\_global\_ids* is not needed by callback functions  
[ZOLTAN\\_PRE\\_MIGRATE\\_PP\\_FN](#), [ZOLTAN\\_MID\\_MIGRATE\\_PP\\_FN](#),  
and [ZOLTAN\\_POST\\_MIGRATE\\_PP\\_FN](#).
- export\_local\_ids* An array of *num\_export* local IDs of objects to be exported to new partitions.  
(size = *num\_export* \* [NUM\\_LID\\_ENTRIES](#))  
All processors may pass *export\_local\_ids*=NULL if import lists are provided and *export\_local\_ids* is not needed by callback functions  
[ZOLTAN\\_PRE\\_MIGRATE\\_PP\\_FN](#), [ZOLTAN\\_MID\\_MIGRATE\\_PP\\_FN](#),  
and [ZOLTAN\\_POST\\_MIGRATE\\_PP\\_FN](#).
- export\_procs* An array of size *num\_export* listing the processor IDs to which exported objects should be assigned (i.e., the destination processors).  
All processors may pass *export\_procs*=NULL if import lists are provided and *export\_procs* is not needed by callback functions  
[ZOLTAN\\_PRE\\_MIGRATE\\_PP\\_FN](#), [ZOLTAN\\_MID\\_MIGRATE\\_PP\\_FN](#),  
and [ZOLTAN\\_POST\\_MIGRATE\\_PP\\_FN](#).

*export\_to\_part* An array of size *num\_export* listing the partitions to which exported objects should be assigned.

All processors may pass *export\_to\_part*=NULL if import lists are provided and *export\_to\_part* is not needed by callback functions

[ZOLTAN\\_PRE\\_MIGRATE\\_PP\\_FN](#), [ZOLTAN\\_MID\\_MIGRATE\\_PP\\_FN](#),  
and [ZOLTAN\\_POST\\_MIGRATE\\_PP\\_FN](#).

### Returned Value:

int [Error code](#).

Note that the number of global and local ID entries ([NUM\\_GID\\_ENTRIES](#) and [NUM\\_LID\\_ENTRIES](#)) should be set using [Zoltan\\_Set\\_Param](#) before calling **Zoltan\_Migrate**. All processors must have the same values for these two parameters.

```
C:      int Zoltan_Compute_Destinations (
        struct Zoltan_Struct *zz,
        int num_known,
        ZOLTAN\_ID\_PTR known_global_ids,
        ZOLTAN\_ID\_PTR known_local_ids,
        int *known_procs,
        int *num_found,
        ZOLTAN\_ID\_PTR *found_global_ids,
        ZOLTAN\_ID\_PTR *found_local_ids,
        int **found_procs);

FORTRAN: FUNCTION Zoltan_Compute_Destinations(zz, num_known, known_global_ids,
        known_local_ids, known_procs, num_found, found_global_ids, found_local_ids,
        found_procs)
        INTEGER(Zoltan_INT) :: Zoltan_Compute_Destinations
        TYPE(Zoltan_Struct),INTENT(IN) :: zz
        INTEGER(Zoltan_INT), INTENT(IN) :: num_known
        INTEGER(Zoltan_INT), INTENT(OUT) :: num_found
        INTEGER(Zoltan_INT), POINTER, DIMENSION(:) :: known_global_ids,
        found_global_ids
        INTEGER(Zoltan_INT), POINTER, DIMENSION(:) :: known_local_ids,
        found_local_ids
        INTEGER(Zoltan_INT), POINTER, DIMENSION(:) :: known_procs, found_procs
```

**Zoltan\_Compute\_Destinations** is a wrapper around [Zoltan\\_Invert\\_Lists](#) that excludes partition assignment arrays. It is maintained for backward compatibility with previous versions of Zoltan.

**Zoltan\_Compute\_Destinations** assumes the number of partitions is equal to the number of processors. The computed lists are allocated in Zoltan; they should not be allocated by the application before calling **Zoltan\_Compute\_Destinations**. These lists can be freed through a call to [Zoltan\\_LB\\_Free\\_Data](#) or [Zoltan\\_LB\\_Free\\_Part](#).



**Arguments:**

All arguments are analogous to those in [Zoltan\\_Invert\\_Lists](#).

Partition-assignment arrays *known\_to\_part* and *found\_to\_part* are not included, as partition and processor numbers are assumed to be the same in **Zoltan\_Compute\_Destinations**.

**Returned Value:**

int            [Error code](#).

Note that the number of global and local ID entries ([NUM\\_GID\\_ENTRIES](#) and [NUM\\_LID\\_ENTRIES](#)) should be set using [Zoltan\\_Set\\_Param](#) before calling **Zoltan\_Compute\_Destinations**. All processors must have the same values for these two parameters.

```
C:      int Zoltan_Help_Migrate (
        struct Zoltan_Struct *zz,
        int num_import,
        ZOLTAN\_ID\_PTR import_global_ids,
        ZOLTAN\_ID\_PTR import_local_ids,
        int *import_procs,
        int num_export,
        ZOLTAN\_ID\_PTR export_global_ids,
        ZOLTAN\_ID\_PTR export_local_ids,
        int *export_procs);

FORTRAN: FUNCTION Zoltan_Help_Migrate(zz, num_import, import_global_ids,
import_local_ids, import_procs, num_export, export_global_ids, export_local_ids,
export_procs)
INTEGER(Zoltan_INT) :: Zoltan_Help_Migrate
TYPE(Zoltan_Struct),INTENT(IN) :: zz
INTEGER(Zoltan_INT), INTENT(IN) :: num_import, num_export
INTEGER(Zoltan_INT), POINTER, DIMENSION(:) :: import_global_ids,
export_global_ids
INTEGER(Zoltan_INT), POINTER, DIMENSION(:) :: import_local_ids,
export_local_ids
INTEGER(Zoltan_INT), POINTER, DIMENSION(:) :: import_procs, export_procs
```

**Zoltan\_Help\_Migrate** is a wrapper around [Zoltan\\_Migrate](#) that excludes partition assignment arrays. It is maintained for backward compatibility with previous versions of Zoltan.

**Zoltan\_Help\_Migrate** assumes the number of partitions is equal to the number of processors. It uses migration pre-, mid-, and post-processing routines [ZOLTAN\\_PRE\\_MIGRATE\\_FN\\_TYPE](#), [ZOLTAN\\_MID\\_MIGRATE\\_FN\\_TYPE](#), and [ZOLTAN\\_POST\\_MIGRATE\\_FN\\_TYPE](#), respectively, which also exclude partition assignment arrays.

**Arguments:**

All arguments are analogous to those in [Zoltan\\_Migrate](#). Partition-assignment arrays *import\_to\_part* and *export\_to\_part* are not included, as partition and processor numbers are assumed to be the same in **Zoltan\_Help\_Migrate**.

**Returned Value:**

int      [Error code](#).

---

[\[Table of Contents\]](#) | [\[Next: Ordering Interface\]](#) | [\[Previous: Functions for Augmenting a Decomposition\]](#)



# Ordering Functions

Zoltan provides limited capability for ordering a set of objects, typically given as a graph. The following functions are the ordering interface functions in the Zoltan library; their descriptions are included below.

## [Zoltan\\_Order](#)

---

```

C:      int Zoltan_Order (
        struct Zoltan_Struct *zz,
        int *num_gid_entries,
        int *num_lid_entries,
        int num_obj,
        ZOLTAN\_ID\_PTR global_ids,
        ZOLTAN\_ID\_PTR local_ids,
        int *rank,
        int *iperm,
        struct Zoltan_Order_Struct *order_info);

FORTRAN: FUNCTION Zoltan_Order(zz, num_gid_entries, num_lid_entries, num_obj,
                                global_ids, local_ids, rank, iperm)
                                INTEGER(Zoltan_INT) :: Zoltan_Order
                                TYPE(Zoltan_Struct), INTENT(IN) :: zz
                                INTEGER(Zoltan_INT), INTENT(OUT) :: num_gid_entries, num_lid_entries
                                INTEGER(Zoltan_INT), INTENT(IN) :: num_obj
                                INTEGER(Zoltan_INT) :: global_ids(*), local_ids(*)
                                INTEGER(Zoltan_INT) :: rank(*), iperm(*)

C++:   int Zoltan::Order (
        int &num_gid_entries,
        int &num_lid_entries,
        const int &num_obj,
        ZOLTAN\_ID\_PTR global_ids,
        ZOLTAN\_ID\_PTR local_ids,
        int *rank,
        int *iperm);

```

**Zoltan\_Order** invokes the ordering routine specified by the [ORDER\\_METHOD](#) parameter. Results of the ordering are returned in the arrays *rank* and *iperm*. *rank[i]* gives the rank of *global\_ids[i]* in the computed ordering, while *iperm* is the inverse permutation of *rank*, that is, *iperm[rank[i]] = i*. The ordering may be either global or local, depending on [ORDER\\_TYPE](#). The arrays *global\_ids*, *local\_ids*, *rank*, and *iperm* should all be allocated by the application before **Zoltan\_Order** is called. Each array must have space for (at least) *num\_obj* elements, where *num\_obj* is the number of objects residing on a processor.

**Arguments:**

<i>zz</i>	Pointer to the Zoltan structure, created by <a href="#">Zoltan_Create</a> , to be used in this invocation of the load-balancing routine.
<i>num_gid_entries</i>	Upon return, the number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	Upon return, the number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>num_obj</i>	Number of objects to order on this processor. At present, <i>num_obj</i> should be the total number of objects residing on a processor. In future releases, ordering only a subset of the objects may be permitted.
<i>global_ids</i>	An array of global IDs of objects to be ordered on this processor. (size = <i>num_obj</i> * <i>num_gid_entries</i> ) The array may be uninitialized on input (if <a href="#">REORDER</a> is false), but memory must have been allocated before <b>Zoltan_Order</b> is called.
<i>local_ids</i>	[Optional.] An array of local IDs of objects to be ordered on this processor. (size = <i>num_obj</i> * <i>num_lid_entries</i> ) The array may be uninitialized on input (if <a href="#">REORDER</a> is false), but memory must have been allocated before <b>Zoltan_Order</b> is called.
<i>rank</i>	Upon return, an array of length <i>num_obj</i> containing the rank of each object in the computed ordering. When <i>rank</i> [ <i>i</i> ] = <i>j</i> , that means that the object corresponding to <i>global_ids</i> [ <i>i</i> ] is the <i>j</i> th object in the ordering. (This array corresponds directly to the <i>perm</i> array in METIS and the <i>order</i> array in ParMETIS.) Note that the rank may refer to either a local or a global ordering, depending on <a href="#">ORDER_TYPE</a> . Memory for this array must have been allocated before <b>Zoltan_Order</b> is called.
<i>iperm</i>	Upon return, an array of length <i>num_obj</i> containing the inverse permutation of <i>rank</i> . That is, <i>iperm</i> [ <i>rank</i> [ <i>i</i> ]] = <i>i</i> . In other words, <i>iperm</i> [ <i>j</i> ] gives the <i>j</i> th object in the ordering. Memory for this array must have been allocated before <b>Zoltan_Order</b> is called.
<i>order_info</i>	Upon return, this struct contains additional information about the ordering produced. This parameter is currently not used and should always be set to NULL. It is not included in the FORTRAN or C++ interface.

**Returned Value:**

int      [Error code](#).

---

[Table of Contents](#) | [Next: Coloring Functions](#) | [Previous: Migration Functions](#)

# Coloring Functions

Zoltan provides limited capability for coloring a set of objects, typically given as a graph. In graph coloring, each vertex is assigned an integer label such that no two adjacent vertices have the same label. The following functions are the coloring interface functions in the Zoltan library; their descriptions are included below.

## [Zoltan\\_Color](#)

---

```
C: int Zoltan_Color (
    struct Zoltan_Struct *zz,
    int *num_gid_entries,
    int *num_lid_entries,
    int num_obj,
    ZOLTAN\_ID\_PTR global_ids,
    ZOLTAN\_ID\_PTR local_ids,
    int *color_exp);
```

---

**Zoltan\_Color** invokes the coloring routine and the assigned colors of each object are returned in the array *color\_exp*. *color\_exp[i]* gives the color of *global\_ids[i]* in the computed coloring. The arrays *global\_ids*, *local\_ids*, and *color\_exp* should all be allocated by the application before **Zoltan\_Color** is called. Each array must have space for (at least) *num\_obj* elements, where *num\_obj* is the number of objects residing on a processor.

### Arguments:

<i>zz</i>	Pointer to the Zoltan structure, created by <a href="#">Zoltan_Create</a> , to be used in this invocation of the load-balancing routine.
<i>num_gid_entries</i>	Upon return, the number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	Upon return, the number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>num_obj</i>	Number of objects to color on this processor. <i>num_obj</i> should be the total number of objects residing on a processor.
<i>global_ids</i>	Upon return, an array of global IDs of objects to be colored on this processor. (size = <i>num_obj</i> * <i>num_gid_entries</i> ) Memory for this array must have been allocated before <b>Zoltan_Color</b> is called.
<i>local_ids</i>	[Optional.] Upon return, an array of local IDs of objects to be colored on this processor. (size = <i>num_obj</i> * <i>num_lid_entries</i> ) Memory for this array must have been allocated before <b>Zoltan_Color</b> is called.

*color\_exp* Upon return, an array of length *num\_obj* containing the colors of objects. That is, *color\_exp[i]* gives the color of *global\_ids[i]* in the computed coloring. (Colors are usually positive integers.) Memory for this array must have been allocated before **Zoltan\_Color** is called.

**Returned Value:**

int      [Error code](#).

---

[Table of Contents](#) | [Next: Application-Registered Query Functions](#) | [Previous: Ordering Functions](#)

# Application-Registered Query Functions

Zoltan gets information about a processor's objects through calls to query functions. These functions must be provided by the application. They are "registered" with Zoltan; that is, a pointer to the function is passed to Zoltan, which can then call that function when its information is needed. Two categories of query functions are used by the library:

[General Zoltan Query Functions](#)

[Migration Query Functions](#)

In each category, a variety of query functions can be registered by the user. The query functions have a function type, describing their purpose. Functions can be registered with a Zoltan structure in two ways: through calls to [Zoltan\\_Set\\_Fn](#) or through calls to query-function-specific functions

[Zoltan\\_Set\\_<zoltan\\_fn\\_type>\\_Fn](#). When a function is registered through a call to [Zoltan\\_Set\\_Fn](#), its function type is passed in the *fn\_type* argument. When [Zoltan\\_Set\\_<zoltan\\_fn\\_type>\\_Fn](#) is used to register functions, the type of the function is implicit in the *fn\_ptr* argument. Each function description below includes both its function type and function prototype.

Query functions that return information about data objects owned by a processor come in two forms: list-based functions that return information about a list of objects, and iterator functions that return information about a single object. Users can provide either version of the query function; they need not provide both. Zoltan calls the list-based functions with the IDs of all objects needed; this approach often provides faster performance as it eliminates the overhead of multiple function calls. List-based functions have the word "MULTI" in their function-type name. If, instead, the application provides iterator functions, Zoltan calls the iterator function once for each object whose data is needed. This approach, while slower, allows Zoltan to use less memory for some data.

Some algorithms in Zoltan require that certain query functions be registered by the application; for example, geometric partitioning algorithms such as Recursive Coordinate Bisection (RCB) require that either a [ZOLTAN\\_GEOM\\_FN](#) or a [ZOLTAN\\_GEOM\\_MULTI\\_FN](#) be registered. When a default value is specified below, the query function type is optional; if a function of that type is not registered, the default values are used. Details of which query functions are required by particular algorithms are included in the [Algorithms](#) section.

Many of the functions have both global and local object identifiers (IDs) in their argument lists. The global IDs provided by the application must be unique across all processors; they are used for identification within Zoltan. The local IDs are not used by Zoltan; they are provided for the convenience of the application and can be anything the application desires. The local IDs can be used by application query routines to enable direct access to application data. For example, the object with global ID "3295" may be stored by the application in location "15" of an array in the processor's local memory. Both global ID "3295" and local ID "15" can be used by the application to describe the object. Then, rather than searching the array for global ID "3295," the application query routines can subsequently use the local ID to index directly into the local storage array. See [Data Types for Object IDs](#) for a description of global and local IDs. All of the functions have, as their first argument, a pointer to data that is passed to

Zoltan through [Zoltan\\_Set\\_Fn](#) or [Zoltan\\_Set\\_<zoltan\\_fn\\_type>\\_Fn](#). This data is not used by Zoltan. A different set of data can be supplied for each registered function. For example, if the local ID is an index into an array of data structures, then the data pointer might point to the head of the data structure array.

As their last argument, all functions have an [error code](#) that should be set and returned by the registered function.

If you are calling the Zoltan library from a C++ application, you may set the query function to be any class static function or any function defined outside of a class definition. However, it is possible you will wish to set the query function to be an object method. In that case, you should write a query function that takes a pointer to the object as its *data* field. The query function can then call the object method.

---

[\[Table of Contents](#) | [Next: Load-Balancing Query Functions](#) | [Previous: Coloring Functions](#)]

# General Zoltan Query Functions

The following registered functions are used by various Zoltan algorithms in the Zoltan library. No single algorithm uses all the query functions; the [algorithm descriptions](#) indicate which query functions are required by individual algorithms.

## [Object ID Functions](#)

[ZOLTAN\\_NUM\\_OBJ\\_FN](#)

[ZOLTAN\\_OBJ\\_LIST\\_FN](#)

[ZOLTAN\\_FIRST\\_OBJ\\_FN](#)

[ZOLTAN\\_NEXT\\_OBJ\\_FN](#)

[ZOLTAN\\_PARTITION\\_MULTI\\_FN](#) or [ZOLTAN\\_PARTITION\\_FN](#)

## [Geometry-Based Functions](#)

[ZOLTAN\\_NUM\\_GEOM\\_FN](#)

[ZOLTAN\\_GEOM\\_MULTI\\_FN](#) or [ZOLTAN\\_GEOM\\_FN](#)

## [Graph-Based Functions](#)

[ZOLTAN\\_NUM\\_EDGES\\_MULTI\\_FN](#) or [ZOLTAN\\_NUM\\_EDGES\\_FN](#)

[ZOLTAN\\_EDGE\\_LIST\\_MULTI\\_FN](#) or [ZOLTAN\\_EDGE\\_LIST\\_FN](#)

## [Hypergraph-Based Functions](#)

[ZOLTAN\\_HG\\_SIZE\\_CS\\_FN](#)

[ZOLTAN\\_HG\\_CS\\_FN](#)

[ZOLTAN\\_HG\\_SIZE\\_EDGE\\_WTS\\_FN](#)

[ZOLTAN\\_HG\\_EDGE\\_WTS\\_FN](#)

## [Tree-Based Functions](#)

[ZOLTAN\\_NUM\\_COARSE\\_OBJ\\_FN](#)

[ZOLTAN\\_COARSE\\_OBJ\\_LIST\\_FN](#)

[ZOLTAN\\_FIRST\\_COARSE\\_OBJ\\_FN](#)

[ZOLTAN\\_NEXT\\_COARSE\\_OBJ\\_FN](#)

[ZOLTAN\\_NUM\\_CHILD\\_FN](#)

[ZOLTAN\\_CHILD\\_LIST\\_FN](#)

[ZOLTAN\\_CHILD\\_WEIGHT\\_FN](#)

## [Border Object Functions \(currently unused\)](#)

[ZOLTAN\\_NUM\\_BORDER\\_OBJ\\_FN](#)

[ZOLTAN\\_BORDER\\_OBJ\\_LIST\\_FN](#)

[ZOLTAN\\_FIRST\\_BORDER\\_OBJ\\_FN](#)

[ZOLTAN\\_NEXT\\_BORDER\\_OBJ\\_FN](#)

---

# Object ID Functions

---

C and C++:   typedef int **ZOLTAN\_NUM\_OBJ\_FN** (void \**data*, int \**ierr*);  
FORTRAN:    FUNCTION *Get\_Num\_Obj*(*data*, *ierr*)  
              INTEGER(Zoltan\_INT) :: Get\_Num\_Obj  
              <*type-data*>, INTENT(IN) :: *data*  
              INTEGER(Zoltan\_INT), INTENT(OUT) :: *ierr*  
  
              <*type-data*> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or  
              REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE),  
              DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section  
              on [Fortran query functions](#) for an explanation.

---

A **ZOLTAN\_NUM\_OBJ\_FN** query function returns the number of objects that are currently assigned to the processor.



**Function Type:** **ZOLTAN\_NUM\_OBJ\_FN\_TYPE**

**Arguments:**

*data* Pointer to user-defined data.  
*ierr* Error code to be set by function.

**Returned Value:**

int The number of objects that are assigned to the processor.

C and C++: `typedef void ZOLTAN_OBJ_LIST_FN (void *data, int num_gid_entries,  
int num_lid_entries, ZOLTAN\_ID\_PTR global_ids, ZOLTAN\_ID\_PTR local_ids,  
int wgt_dim, float *obj_wgts, int *ierr);`

FORTTRAN: `SUBROUTINE Get_Obj_List(data, num_gid_entries, num_lid_entries, global_ids,  
local_ids, wgt_dim, obj_wgts, ierr)  
<type-data>, INTENT(IN) :: data  
INTEGER(Zoltan_INT), INTENT(IN) :: num_gid_entries, num_lid_entries  
INTEGER(Zoltan_INT), INTENT(OUT), DIMENSION(*) :: global_ids  
INTEGER(Zoltan_INT), INTENT(OUT), DIMENSION(*) :: local_ids  
INTEGER(Zoltan_INT), INTENT(IN) :: wgt_dim  
REAL(Zoltan_FLOAT), INTENT(OUT), DIMENSION(*) :: obj_wgts  
INTEGER(Zoltan_INT), INTENT(OUT) :: ierr`

<type-data> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or  
REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE),  
DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section  
on [Fortran query functions](#) for an explanation.

A **ZOLTAN\_OBJ\_LIST\_FN** query function fills two (three if weights are used) arrays with information about the objects currently assigned to the processor. Both arrays are allocated (and subsequently freed) by Zoltan; their size is determined by a call to a [ZOLTAN\\_NUM\\_OBJ\\_FN](#) query function to get the array size. For many algorithms, either a **ZOLTAN\_OBJ\_LIST\_FN** query function or a [ZOLTAN\\_FIRST\\_OBJ\\_FN/ZOLTAN\\_NEXT\\_OBJ\\_FN](#) query-function pair must be registered; however, both query options need not be provided.

**Function Type:** **ZOLTAN\_OBJ\_LIST\_FN\_TYPE**

**Arguments:**

*data* Pointer to user-defined data.  
*num\_gid\_entries* The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter [NUM\\_GID\\_ENTRIES](#).  
*num\_lid\_entries* The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter [NUM\\_LID\\_ENTRIES](#).  
*global\_ids* Upon return, an array of unique global IDs for all objects assigned to the processor.  
*local\_ids* Upon return, an array of local IDs, the meaning of which can be determined by the application, for all objects assigned to the processor.  
*wgt\_dim* The number of weights associated with an object (typically 1), or 0 if weights are not requested. This value is set through the parameter [OBJ\\_WEIGHT\\_DIM](#).  
*obj\_wgts* Upon return, an array of object weights. Weights for object *i* are stored in *obj\_wgts*[(*i*-1)\**wgt\_dim*:*i*\**wgt\_dim*-1]. If *wgt\_dim*=0, the return value of *obj\_wgts* is undefined and may be NULL.  
*ierr* Error code to be set by function.

C and C++: `typedef int ZOLTAN_FIRST_OBJ_FN (void *data, int num_gid_entries,  
int num_lid_entries, ZOLTAN\_ID\_PTR first_global_id,  
ZOLTAN\_ID\_PTR first_local_id, int wgt_dim, float *first_obj_wgt, int *ierr);`



**FORTRAN:**    **FUNCTION** *Get\_First\_Obj*(*data*, *num\_gid\_entries*, *num\_lid\_entries*, *first\_global\_id*,  
*first\_local\_id*, *wgt\_dim*, *first\_obj\_wgt*, *ierr*)  
 INTEGER(Zoltan\_INT) :: Get\_First\_Obj  
 <type-data>, INTENT(IN) :: *data*  
 INTEGER(Zoltan\_INT), INTENT(IN) :: *num\_gid\_entries*, *num\_lid\_entries*  
 INTEGER(Zoltan\_INT), INTENT(OUT), DIMENSION(\*) :: *first\_global\_id*  
 INTEGER(Zoltan\_INT), INTENT(OUT), DIMENSION(\*) :: *first\_local\_id*  
 INTEGER(Zoltan\_INT), INTENT(IN) :: *wgt\_dim*  
 REAL(Zoltan\_FLOAT), INTENT(OUT), DIMENSION(\*) :: *first\_obj\_wgt*  
 INTEGER(Zoltan\_INT), INTENT(OUT) :: *ierr*

<type-data> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or  
 REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE),  
 DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section  
 on [Fortran query functions](#) for an explanation.

---

A **ZOLTAN\_FIRST\_OBJ\_FN** query function initializes an iteration over objects assigned to the processor. It returns the global and local IDs of the first object on the processor. Subsequent calls to a [ZOLTAN\\_NEXT\\_OBJ\\_FN](#) query function iterate over and return other objects assigned to the processor. This query-function pair frees the application from having to build an array of objects (as in [ZOLTAN\\_OBJ\\_LIST\\_FN](#)) and allows Zoltan's routines to obtain only as much information about objects as they need. For many algorithms, either a [ZOLTAN\\_OBJ\\_LIST\\_FN](#) query function or a **ZOLTAN\_FIRST\_OBJ\_FN**/[ZOLTAN\\_NEXT\\_OBJ\\_FN](#) query-function pair must be registered; however, both query options need not be provided.

**Function Type:**    **ZOLTAN\_FIRST\_OBJ\_FN\_TYPE**

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>first_global_id</i>	The returned value of the global ID for the first object; the value is ignored if there are no objects.
<i>first_local_id</i>	The returned value of the local ID for the first object; the value is ignored if there are no objects.
<i>wgt_dim</i>	The number of weights associated with an object (typically 1), or 0 if weights are not requested. This value is set through the parameter <a href="#">OBJ_WEIGHT_DIM</a> .
<i>first_obj_wgt</i>	Upon return, the first object's weights; an array of length <i>wgt_dim</i> . Undefined if <i>wgt_dim</i> =0.
<i>ierr</i>	Error code to be set by function.

**Returned Value:**

1	If <i>first_global_id</i> and <i>first_local_id</i> contain valid IDs of the first object.
0	If no objects are available.

---

**C and C++:**    typedef int **ZOLTAN\_NEXT\_OBJ\_FN** (void \* *data*, int *num\_gid\_entries*,  
 int *num\_lid\_entries*, [ZOLTAN\\_ID\\_PTR](#) *global\_id*, [ZOLTAN\\_ID\\_PTR](#) *local\_id*,  
[ZOLTAN\\_ID\\_PTR](#) *next\_global\_id*, [ZOLTAN\\_ID\\_PTR](#) *next\_local\_id*, int *wgt\_dim*,  
 float \**next\_obj\_wgt*, int \**ierr*);

**FORTRAN:**    **FUNCTION** *Get\_Next\_Obj*(*data*, *num\_gid\_entries*, *num\_lid\_entries*, *global\_id*,  
*local\_id*, *next\_global\_id*, *next\_local\_id*, *wgt\_dim*, *next\_obj\_wgt*, *ierr*)  
 INTEGER(Zoltan\_INT) :: Get\_Next\_Obj  
 <type-data>, INTENT(IN) :: *data*  
 INTEGER(Zoltan\_INT), INTENT(IN) :: *num\_gid\_entries*, *num\_lid\_entries*  
 INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *global\_id*  
 INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *local\_id*  
 INTEGER(Zoltan\_INT), INTENT(OUT), DIMENSION(\*) :: *next\_global\_id*  
 INTEGER(Zoltan\_INT), INTENT(OUT), DIMENSION(\*) :: *next\_local\_id*  
 INTEGER(Zoltan\_INT), INTENT(IN) :: *wgt\_dim*  
 REAL(Zoltan\_FLOAT), INTENT(OUT), DIMENSION(\*) :: *next\_obj\_wgt*  
 INTEGER(Zoltan\_INT), INTENT(OUT) :: *ierr*

*<type-data>* can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE), DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

A **ZOLTAN\_NEXT\_OBJ\_FN** query function is an iterator function which, when given an object assigned to the processor, returns the next object assigned to the processor. The first object of the iteration is provided by a [ZOLTAN\\_FIRST\\_OBJ\\_FN](#) query function. This query-function pair frees the application from having to build an array of objects (as in [ZOLTAN\\_OBJ\\_LIST\\_FN](#)) and allows Zoltan's routines to obtain only as much information about objects as they need. For many algorithms, either a [ZOLTAN\\_OBJ\\_LIST\\_FN](#) query function or a [ZOLTAN\\_FIRST\\_OBJ\\_FN](#)/ZOLTAN\_NEXT\_OBJ\_FN query-function pair must be registered; however, both query options need not be provided.

**Function Type:** ZOLTAN\_NEXT\_OBJ\_FN\_TYPE

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>global_id</i>	The global ID of the previous object.
<i>local_id</i>	The local ID of the previous object.
<i>next_global_id</i>	The returned value of the global ID for the next object; the value is ignored if there are no more objects.
<i>next_local_id</i>	The returned value of the local ID for the next object; the value is ignored if there are no more objects.
<i>wgt_dim</i>	The number of weights associated with an object (typically 1), or 0 if weights are not requested. This value is set through the parameter <a href="#">OBJ_WEIGHT_DIM</a> .
<i>next_obj_wgt</i>	Upon return, the next object's weights; an array of length <i>wgt_dim</i> . Undefined if <i>wgt_dim</i> =0.
<i>ierr</i>	Error code to be set by function.

**Returned Value:**

1	If <i>next_global_id</i> and <i>next_local_id</i> contain valid IDs of the next object.
0	If no more objects are available.

C and C++: typedef void **ZOLTAN\_PARTITION\_MULTI\_FN** (void \**data*, int *num\_gid\_entries*, int *num\_lid\_entries*, int *num\_obj*, [ZOLTAN\\_ID\\_PTR](#) *global\_ids*, [ZOLTAN\\_ID\\_PTR](#) *local\_ids*, int \**parts*, int \**ierr*);

FORTRAN: SUBROUTINE *Get\_Partition\_Multi*(*data*, *num\_gid\_entries*, *num\_lid\_entries*, *num\_obj*, *global\_ids*, *local\_ids*, *ierr*)  
*<type-data>*, INTENT(IN) :: *data*  
 INTEGER(Zoltan\_INT), INTENT(IN) :: *num\_gid\_entries*, *num\_lid\_entries*, *num\_obj*  
 INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *global\_ids*  
 INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *local\_ids*  
 INTEGER(Zoltan\_INT), INTENT(OUT), DIMENSION(\*) :: *parts*  
 INTEGER(Zoltan\_INT), INTENT(OUT) :: *ierr*

*<type-data>* can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE), DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

A **ZOLTAN\_PARTITION\_MULTI\_FN** query function returns a list of partitions to which given objects are currently assigned. If a **ZOLTAN\_PARTITION\_MULTI\_FN** or [ZOLTAN\\_PARTITION\\_FN](#) is not registered, Zoltan assumes the partition numbers are the processor number of the owning processor. Valid partition numbers are non-negative integers.

**Function Type:** ZOLTAN\_PARTITION\_MULTI\_FN\_TYPE

**Arguments:**

<i>data</i>	Pointer to user-defined data.
-------------	-------------------------------

<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>num_obj</i>	The number of object IDs in arrays <i>global_ids</i> and <i>local_ids</i> .
<i>global_ids</i>	The global IDs of the objects for which the partition numbers should be returned.
<i>local_ids</i>	The local IDs of the objects for which the partition numbers should be returned.
<i>parts</i>	Upon return, an array of partition numbers corresponding to the global and local IDs.
<i>ierr</i>	Error code to be set by function.

```
C and C++:  typedef int ZOLTAN_PARTITION_FN (void *data, int num_gid_entries,
int num_lid_entries, ZOLTAN\_ID\_PTR global_id, ZOLTAN\_ID\_PTR local_id,
int *ierr);

FORTRAN:   FUNCTION Get_Partition(data, num_gid_entries, num_lid_entries, global_id,
local_id, ierr)
INTEGER(Zoltan_INT) :: Get_Partition
<type-data>, INTENT(IN) :: data
INTEGER(Zoltan_INT), INTENT(IN) :: num_gid_entries, num_lid_entries
INTEGER(Zoltan_INT), INTENT(IN), DIMENSION(*) :: global_id
INTEGER(Zoltan_INT), INTENT(IN), DIMENSION(*) :: local_id
INTEGER(Zoltan_INT), INTENT(OUT) :: ierr

<type-data> can be any of INTEGER(Zoltan_INT), DIMENSION(*) or
REAL(Zoltan_FLOAT), DIMENSION(*) or REAL(Zoltan_DOUBLE),
DIMENSION(*) or TYPE(Zoltan_User_Data_x) where x is 1, 2, 3 or 4. See the section
on Fortran query functions for an explanation.
```

A **ZOLTAN\_PARTITION\_FN** query function returns the partition to which a given object is currently assigned. If a **ZOLTAN\_PARTITION\_FN** or [ZOLTAN\\_PARTITION\\_MULTI\\_FN](#) is not registered, Zoltan assumes the partition numbers are the processor number of the owning processor. Valid partition numbers are non-negative integers.

**Function Type:**     **ZOLTAN\_PARTITION\_FN\_TYPE**

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>global_id</i>	The global ID of the object for which the partition number should be returned.
<i>local_id</i>	The local ID of the object for which the partition number should be returned.
<i>ierr</i>	Error code to be set by function.

**Returned Value:**

int	The partition number for the object identified by <i>global_id</i> and <i>local_id</i> .
-----	--

# Geometry-based Functions

```
C and C++:  typedef int ZOLTAN_NUM_GEOM_FN (void *data, int *ierr);
```

**FORTRAN:**     **FUNCTION *Get\_Num\_Geom*(data, ierr)**  
                   **INTEGER(Zoltan\_INT) :: *Get\_Num\_Geom***  
                   <*type-data*>, **INTENT(IN) :: data**  
                   **INTEGER(Zoltan\_INT), INTENT(OUT) :: ierr**

<*type-data*> can be any of **INTEGER(Zoltan\_INT)**, **DIMENSION(\*)** or **REAL(Zoltan\_FLOAT)**, **DIMENSION(\*)** or **REAL(Zoltan\_DOUBLE)**, **DIMENSION(\*)** or **TYPE(Zoltan\_User\_Data\_x)** where *x* is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

A **ZOLTAN\_NUM\_GEOM\_FN** query function returns the number of values needed to express the geometry of an object. For example, for a two-dimensional mesh-based application, (x,y) coordinates are needed to describe an object's geometry; thus the **ZOLTAN\_NUM\_GEOM\_FN** query function should return the value of two. For a similar three-dimensional application, the return value should be three.

**Function Type:**     **ZOLTAN\_NUM\_GEOM\_FN\_TYPE**

**Arguments:**

*data*                 Pointer to user-defined data.  
*ierr*                 Error code to be set by function.

**Returned Value:**

int                   The number of values needed to express the geometry of an object.

**C and C++:**     **typedef void **ZOLTAN\_GEOM\_MULTI\_FN**** (void \**data*, int *num\_gid\_entries*,  
                   int *num\_lid\_entries*, int *num\_obj*, [ZOLTAN\\_ID\\_PTR](#) *global\_ids*,  
                   [ZOLTAN\\_ID\\_PTR](#) *local\_ids*, int *num\_dim*, double \**geom\_vec*, int \**ierr*);

**FORTRAN:**     **SUBROUTINE *Get\_Geom\_Multi*(data, num\_gid\_entries, num\_lid\_entries, num\_obj,**  
                   *global\_ids*, *local\_ids*, *num\_dim*, *geom\_vec*, *ierr*)  
                   <*type-data*>, **INTENT(IN) :: data**  
                   **INTEGER(Zoltan\_INT), INTENT(IN) :: num\_gid\_entries, num\_lid\_entries**  
                   **INTEGER(Zoltan\_INT), INTENT(IN) :: num\_obj, num\_dim**  
                   **INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: global\_ids**  
                   **INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: local\_ids**  
                   **REAL(Zoltan\_DOUBLE), INTENT(OUT), DIMENSION(\*) :: geom\_vec**  
                   **INTEGER(Zoltan\_INT), INTENT(OUT) :: ierr**

<*type-data*> can be any of **INTEGER(Zoltan\_INT)**, **DIMENSION(\*)** or **REAL(Zoltan\_FLOAT)**, **DIMENSION(\*)** or **REAL(Zoltan\_DOUBLE)**, **DIMENSION(\*)** or **TYPE(Zoltan\_User\_Data\_x)** where *x* is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

A **ZOLTAN\_GEOM\_MULTI\_FN** query function returns a vector of geometry values for a list of given objects. The geometry vector is allocated by Zoltan to be of size *num\_obj* \* *num\_dim*; its format is described below.

**Function Type:**     **ZOLTAN\_GEOM\_MULTI\_FN\_TYPE**

**Arguments:**

*data*                 Pointer to user-defined data.  
*num\_gid\_entries*     The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter [NUM\\_GID\\_ENTRIES](#).  
*num\_lid\_entries*     The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter [NUM\\_LID\\_ENTRIES](#).  
*num\_obj*             The number of object IDs in arrays *global\_ids* and *local\_ids*.  
*global\_ids*           Array of global IDs of objects whose geometry values should be returned.  
*local\_ids*            Array of local IDs of objects whose geometry values should be returned.  
*num\_dim*             Number of coordinate entries per object (typically 1, 2, or 3).  
*geom\_vec*            Upon return, an array containing geometry values. For object *i* (specified by *global\_ids*[*i*\**num\_gid\_entries*] and *local\_ids*[*i*\**num\_lid\_entries*], *i*=0,1,...,*num\_obj*-1), coordinate values should be stored in *geom\_vec*[*i*\**num\_dim*:(*i*+1)\**num\_dim*-1].  
*ierr*                 Error code to be set by function.

---

C and C++:    typedef void **ZOLTAN\_GEOM\_FN** (void \**data*, int *num\_gid\_entries*,  
                   int *num\_lid\_entries*, [ZOLTAN\\_ID\\_PTR](#) *global\_id*, [ZOLTAN\\_ID\\_PTR](#) *local\_id*,  
                   double \**geom\_vec*, int \**ierr*);

FORTRAN:    SUBROUTINE *Get\_Geom*(*data*, *num\_gid\_entries*, *num\_lid\_entries*, *global\_id*,  
                   *local\_id*, *geom\_vec*, *ierr*)  
               <*type-data*>, INTENT(IN) :: *data*  
               INTEGER(Zoltan\_INT), INTENT(IN) :: *num\_gid\_entries*, *num\_lid\_entries*  
               INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *global\_id*  
               INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *local\_id*  
               REAL(Zoltan\_DOUBLE), INTENT(OUT), DIMENSION(\*) :: *geom\_vec*  
               INTEGER(Zoltan\_INT), INTENT(OUT) :: *ierr*

              <*type-data*> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or  
               REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE),  
               DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_*x*) where *x* is 1, 2, 3 or 4. See the section  
               on [Fortran query functions](#) for an explanation.

---

A **ZOLTAN\_GEOM\_FN** query function returns a vector of geometry values for a given object. The geometry vector is allocated by Zoltan to be of the size returned by a [ZOLTAN\\_NUM\\_GEOM\\_FN](#) query function.

**Function Type:**    **ZOLTAN\_GEOM\_FN\_TYPE**

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>global_id</i>	The global ID of the object whose geometry values should be returned.
<i>local_id</i>	The local ID of the object whose geometry values should be returned.
<i>geom_vec</i>	Upon return, an array containing geometry values.
<i>ierr</i>	Error code to be set by function.

---

## Graph-based Functions

---

C and C++:    typedef void **ZOLTAN\_NUM\_EDGES\_MULTI\_FN** (void \**data*,  
                   int *num\_gid\_entries*, int *num\_lid\_entries*, int *num\_obj*, [ZOLTAN\\_ID\\_PTR](#) *global\_ids*,  
                   [ZOLTAN\\_ID\\_PTR](#) *local\_ids*, int \**num\_edges*, int \**ierr*);

FORTRAN:    SUBROUTINE *Get\_Num\_Edges\_Multi*(*data*, *num\_gid\_entries*, *num\_lid\_entries*,  
                   *num\_obj*, *global\_ids*, *local\_ids*, *num\_edges*, *ierr*)  
               INTEGER(Zoltan\_INT) :: *Get\_Num\_Edges*  
               <*type-data*>, INTENT(IN) :: *data*  
               INTEGER(Zoltan\_INT), INTENT(IN) :: *num\_gid\_entries*, *num\_lid\_entries*, *num\_obj*  
               INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *global\_ids*  
               INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *local\_ids*  
               INTEGER(Zoltan\_INT), INTENT(OUT), DIMENSION(\*) :: *num\_edges*  
               INTEGER(Zoltan\_INT), INTENT(OUT) :: *ierr*

              <*type-data*> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or  
               REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE),  
               DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_*x*) where *x* is 1, 2, 3 or 4. See the section  
               on [Fortran query functions](#) for an explanation.

---

A **ZOLTAN\_NUM\_EDGES\_MULTI\_FN** query function returns the number of edges in the communication graph of the application for each



object in a list of objects. That is, for each object in the *global\_ids/local\_ids* arrays, the number of objects with which the given object must share information is returned.

**Function Type:** **ZOLTAN\_NUM\_EDGES\_MULTI\_FN\_TYPE**

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>num_obj</i>	The number of object IDs in arrays <i>global_ids</i> and <i>local_ids</i> .
<i>global_ids</i>	Array of global IDs of objects whose number of edges should be returned.
<i>local_ids</i>	Array of local IDs of objects whose number of edges should be returned.
<i>num_edges</i>	Upon return, an array containing numbers of edges. For object <i>i</i> (specified by <i>global_ids[i*num_gid_entries]</i> and <i>local_ids[i*num_lid_entries]</i> , <i>i=0,1,...,num_obj-1</i> ), the number of edges should be stored in <i>num_edges[i]</i> .
<i>ierr</i>	Error code to be set by function.

C and C++: `typedef int ZOLTAN_NUM_EDGES_FN (void *data, int num_gid_entries, int num_lid_entries, ZOLTAN\_ID\_PTR global_id, ZOLTAN\_ID\_PTR local_id, int *ierr);`

FORTRAN: `FUNCTION Get_Num_Edges(data, num_gid_entries, num_lid_entries, global_id, local_id, ierr)  
 INTEGER(Zoltan_INT) :: Get_Num_Edges  
 <type-data>, INTENT(IN) :: data  
 INTEGER(Zoltan_INT), INTENT(IN) :: num_gid_entries, num_lid_entries  
 INTEGER(Zoltan_INT), INTENT(IN), DIMENSION(*) :: global_id  
 INTEGER(Zoltan_INT), INTENT(IN), DIMENSION(*) :: local_id  
 INTEGER(Zoltan_INT), INTENT(OUT) :: ierr`

<type-data> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE), DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

A **ZOLTAN\_NUM\_EDGES\_FN** query function returns the number of edges for a given object in the communication graph of the application (i.e., the number of objects with which the given object must share information).

**Function Type:** **ZOLTAN\_NUM\_EDGES\_FN\_TYPE**

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>global_id</i>	The global ID of the object for which the number of edges should be returned.
<i>local_id</i>	The local ID of the object for which the number of edges should be returned.
<i>ierr</i>	Error code to be set by function.

**Returned Value:**

int	The number of edges for the object identified by <i>global_id</i> and <i>local_id</i> .
-----	---

C and C++: `typedef void ZOLTAN_EDGE_LIST_MULTI_FN (void *data, int num_gid_entries, int num_lid_entries, int num_obj, ZOLTAN\_ID\_PTR global_ids, ZOLTAN\_ID\_PTR local_ids, int *num_edges, ZOLTAN\_ID\_PTR nbor_global_id, int *nbor_procs, int wgt_dim, float *ewgts, int *ierr);`

**FORTRAN:** SUBROUTINE *Get\_Edge\_List\_Multi*(*data*, *num\_gid\_entries*, *num\_lid\_entries*,  
*num\_obj*, *global\_ids*, *local\_ids*, *num\_edges*, *nbor\_global\_id*, *nbor\_procs*, *wgt\_dim*,  
*ewgts*, *ierr*)  
<type-data>, INTENT(IN) :: *data*  
INTEGER(Zoltan\_INT), INTENT(IN) :: *num\_gid\_entries*, *num\_lid\_entries*, *num\_obj*  
INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *global\_ids*  
INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *local\_ids*  
INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *num\_edges*  
INTEGER(Zoltan\_INT), INTENT(OUT), DIMENSION(\*) :: *nbor\_global\_id*  
INTEGER(Zoltan\_INT), INTENT(OUT), DIMENSION(\*) :: *nbor\_procs*  
INTEGER(Zoltan\_INT), INTENT(IN) :: *wgt\_dim*  
REAL(Zoltan\_FLOAT), INTENT(OUT), DIMENSION(\*) :: *ewgts*  
INTEGER(Zoltan\_INT), INTENT(OUT) :: *ierr*

<type-data> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or  
REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE),  
DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section  
on [Fortran query functions](#) for an explanation.

A **ZOLTAN\_EDGE\_LIST\_MULTI\_FN** query function returns lists of global IDs, processor IDs, and optionally edge weights for objects sharing edges with objects specified in the *global\_ids* input array; objects share edges when they must share information with other objects. The arrays for the returned neighbor lists are allocated by Zoltan; their size is determined by a calls to [ZOLTAN\\_NUM\\_EDGES\\_MULTI\\_FN](#) or [ZOLTAN\\_NUM\\_EDGES\\_FN](#) query functions.

**Function Type:** ZOLTAN\_EDGE\_LIST\_MULTI\_FN\_TYPE

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>num_obj</i>	The number of object IDs in arrays <i>global_ids</i> and <i>local_ids</i> .
<i>global_ids</i>	Array of global IDs of objects whose edge lists should be returned.
<i>local_ids</i>	Array of local IDs of objects whose edge lists should be returned.
<i>num_edges</i>	An array containing numbers of edges for each object in <i>global_ids</i> . For object <i>i</i> (specified by <i>global_ids</i> [ <i>i</i> * <i>num_gid_entries</i> ] and <i>local_ids</i> [ <i>i</i> * <i>num_lid_entries</i> ], <i>i</i> =0,1,..., <i>num_obj</i> -1), the number of edges is stored in <i>num_edges</i> [ <i>i</i> ].
<i>nbor_global_id</i>	Upon return, an array of global IDs of objects sharing edges with the objects specified in <i>global_ids</i> . For object <i>i</i> (specified by <i>global_ids</i> [ <i>i</i> * <i>num_gid_entries</i> ] and <i>local_ids</i> [ <i>i</i> * <i>num_lid_entries</i> ], <i>i</i> =0,1,..., <i>num_obj</i> -1), edges are stored in <i>nbor_global_id</i> [ <i>sum</i> * <i>num_gid_entries</i> ] to <i>nbor_global_id</i> [( <i>sum</i> + <i>num_edges</i> [ <i>i</i> ])* <i>num_gid_entries</i> -1], where <i>sum</i> = the sum of <i>num_edges</i> [ <i>j</i> ] for <i>j</i> =0,1,..., <i>i</i> -1.
<i>nbor_procs</i>	Upon return, an array of processor IDs that identifies where the neighboring objects reside. For neighboring object <i>i</i> (stored in <i>nbor_global_id</i> [ <i>i</i> * <i>num_gid_entries</i> ]), the processor owning the neighbor is stored in <i>nbor_procs</i> [ <i>i</i> ].
<i>wgt_dim</i>	The number of weights associated with an edge (typically 1), or 0 if edge weights are not requested. This value is set through the parameter <a href="#">EDGE_WEIGHT_DIM</a> .
<i>ewgts</i>	Upon return, an array of edge weights, where <i>ewgts</i> [ <i>i</i> * <i>wgt_dim</i> :( <i>i</i> +1)* <i>wgt_dim</i> -1] corresponds to the weights for the <i>i</i> th edge. If <i>wgt_dim</i> =0, the return value of <i>ewgts</i> is undefined and may be NULL.
<i>ierr</i>	Error code to be set by function.

**C and C++:** typedef void **ZOLTAN\_EDGE\_LIST\_FN** (void \**data*, int *num\_gid\_entries*,  
int *num\_lid\_entries*, [ZOLTAN\\_ID\\_PTR](#) *global\_id*, [ZOLTAN\\_ID\\_PTR](#) *local\_id*,  
[ZOLTAN\\_ID\\_PTR](#) *nbor\_global\_id*, int \**nbor\_procs*, int *wgt\_dim*, float \**ewgts*,  
int \**ierr*);

**FORTTRAN:** SUBROUTINE *Get\_Edge\_List*(*data*, *num\_gid\_entries*, *num\_lid\_entries*, *global\_id*, *local\_id*, *nbor\_global\_id*, *nbor\_procs*, *wgt\_dim*, *ewgts*, *ierr*)  
 <type-data>, INTENT(IN) :: *data*  
 INTEGER(Zoltan\_INT), INTENT(IN) :: *num\_gid\_entries*, *num\_lid\_entries*  
 INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *global\_id*  
 INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *local\_id*  
 INTEGER(Zoltan\_INT), INTENT(OUT), DIMENSION(\*) :: *nbor\_global\_id*  
 INTEGER(Zoltan\_INT), INTENT(OUT), DIMENSION(\*) :: *nbor\_procs*  
 INTEGER(Zoltan\_INT), INTENT(IN) :: *wgt\_dim*  
 REAL(Zoltan\_FLOAT), INTENT(OUT), DIMENSION(\*) :: *ewgts*  
 INTEGER(Zoltan\_INT), INTENT(OUT) :: *ierr*

<type-data> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE), DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

---

A **ZOLTAN\_EDGE\_LIST\_FN** query function returns lists of global IDs, processor IDs, and optionally edge weights for objects sharing an edge with a given object (i.e., objects that must share information with the given object). The arrays for the returned neighbor lists are allocated by Zoltan; their size is determined by a call to [ZOLTAN\\_NUM\\_EDGES\\_MULTI\\_FN](#) or [ZOLTAN\\_NUM\\_EDGES\\_FN](#) query functions.

**Function Type:** ZOLTAN\_EDGE\_LIST\_FN\_TYPE

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>global_id</i>	The global ID of the object for which an edge list should be returned.
<i>local_id</i>	The local ID of the object for which an edge list should be returned.
<i>nbor_global_id</i>	Upon return, an array of global IDs of objects sharing edges with the given object.
<i>nbor_procs</i>	Upon return, an array of processor IDs that identifies where the neighboring objects reside.
<i>wgt_dim</i>	The number of weights associated with an edge (typically 1), or 0 if edge weights are not requested. This value is set through the parameter <a href="#">EDGE_WEIGHT_DIM</a> .
<i>ewgts</i>	Upon return, an array of edge weights, where <i>ewgts</i> [ <i>i</i> * <i>wgt_dim</i> :( <i>i</i> +1)* <i>wgt_dim</i> -1] corresponds to the weights for the <i>i</i> th edge. If <i>wgt_dim</i> =0, the return value of <i>ewgts</i> is undefined and may be NULL.
<i>ierr</i>	Error code to be set by function.

---

## Hypergraph-based Functions

---

**C and C++:** typedef void **ZOLTAN\_HG\_SIZE\_CS\_FN** (void \**data*, int \**num\_lists*, int \**num\_pins*, int \**format*, int \**ierr*);

**FORTTRAN:** SUBROUTINE *Get\_HG\_Size\_CS*(*data*, *num\_lists*, *num\_pins*, *format*, *ierr*)  
 <type-data>, INTENT(IN) :: *data*  
 INTEGER(Zoltan\_INT), INTENT(OUT) :: *num\_lists*  
 INTEGER(Zoltan\_INT), INTENT(OUT) :: *num\_pins*  
 INTEGER(Zoltan\_INT), INTENT(OUT) :: *format*  
 INTEGER(Zoltan\_INT), INTENT(OUT) :: *ierr*

<type-data> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE), DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.



A hypergraph (which may alternatively be viewed as a sparse matrix) can be supplied to the Zoltan library in one of two compressed storage formats. In compressed hyperedge format (**ZOLTAN\_COMPRESSED\_EDGE**) a list of global hyperedge IDs is provided. Then a single list of the hypergraph pins, is provided. A pin is the connection between a vertex and a hyperedge (corresponds to a nonzero in a sparse matrix). Pins do not have separate IDs but are rather identified by the global ID of the vertex containing the pin, and implicitly also by the hyperedge ID. An example is provided below.

The other format is compressed vertex (**ZOLTAN\_COMPRESSED\_VERTEX**). In this format a list of vertex global IDs is provided. Then a list of pins ordered by vertex and then by hyperedge is provided. The pin ID in this case is the global ID of the row (or hyperedge) in which the pin appears. In both formats, an array must be provided pointing to the start in the list of pins where each row or column begins. Sparse matrix users may think of these two formats as CSR (compressed sparse row) and CSC (compressed sparse column) format, respectively.

The point of this query function is to tell Zoltan in which format the application will supply the hypergraph, how many vertices and hyperedges there will be, and how many pins. The actual hypergraph is supplied with a query function of the type [\*\*ZOLTAN\\_HG\\_CS\\_FN\\_TYPE\*\*](#).

This query function is required by all applications using the hypergraph methods of Zoltan (unless they are using the [\*\*graph-based\*\*](#) functions with hypergraph code instead).

**Function Type:** **ZOLTAN\_HG\_SIZE\_CS\_FN\_TYPE**

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_lists</i>	Upon return, the number of rows (if using compressed row storage) or columns (if using compressed column storage) that will be supplied to Zoltan by the application process.
<i>num_pins</i>	Upon return, the number of pins (matrix non-zeroes) that will be supplied to Zoltan by the application process.
<i>format</i>	Upon return, the format in which the application process will provide the hypergraph to Zoltan. The options are <b>ZOLTAN_COMPRESSED_EDGE</b> and <b>ZOLTAN_COMPRESSED_VERTEX</b> .
<i>ierr</i>	Error code to be set by function.

---

C and C++: `typedef void ZOLTAN_HG_CS_FN (void *data, int num_gid_entries,  
int num_row_or_col, int num_pins, int format, ZOLTAN\_ID\_PTR vtxedge_GID,  
int *vtxedge_ptr, ZOLTAN\_ID\_PTR pin_GID, int *ierr);`

FORTTRAN: `SUBROUTINE Get_HG_CS(data, num_gid_entries, num_row_or_col, num_pins,  
format, vtxedge_GID, vtxedge_ptr, pin_GID, ierr)  
<type-data>, INTENT(IN) :: data  
INTEGER(Zoltan_INT), INTENT(IN) :: num_gid_entries, num_row_or_col, num_pins,  
format  
INTEGER(Zoltan_INT), INTENT(OUT), DIMENSION(*) :: vtxedge_GID  
INTEGER(Zoltan_INT), INTENT(OUT), DIMENSION(*) :: vtxedge_ptr  
INTEGER(Zoltan_INT), INTENT(OUT), DIMENSION(*) :: pin_GID  
INTEGER(Zoltan_INT), INTENT(OUT) :: ierr`

<type-data> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE), DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

---

**A `ZOLTAN_HG_CS_FN` returns a hypergraph, in a sparse matrix-like style. The size and format of the data to be returned must have been supplied to Zoltan using a [ZOLTAN\\_HG\\_SIZE\\_CS\\_FN\\_TYPE](#) function.**

**When a hypergraph is distributed across multiple processes, Zoltan expects that all processes share a consistent global numbering scheme for hyperedges and vertices. Also, no two processes should return the same pin (matrix non-zero) in this query function.**

**This query function is required by all applications using the hypergraph methods of Zoltan (unless they are using the [graph-based](#) functions with hypergraph code instead).**

**Function Type:** `ZOLTAN_HG_CS_FN_TYPE`

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_vtx_edge</i>	The number of global IDs that is expected to appear on return in <i>vtxedge_GID</i> .
<i>num_pins</i>	The number of pins that is expected to appear on return in <i>pin_GID</i> .
<i>format</i>	If <i>format</i> is <code>ZOLTAN_COMPRESSED_EDGE</code> , Zoltan expects that row (hyperedge) global IDs will be returned in <i>vtxedge_GID</i> , and that column (vertex) global IDs will be returned in <i>pin_GIDs</i> . If it is <code>ZOLTAN_COMPRESSED_VERTEX</code> , then column global IDs are expected to be returned in <i>vtxedge_GID</i> and row global IDs are expected to be returned in <i>pin_GIDs</i> .
<i>vtxedge_GID</i>	Upon return, a list of <i>num_row_or_col</i> global IDs.
<i>vtxedge_ptr</i>	Upon return, this array contains <i>num_row_or_col</i> integers. The integer in the <i>i</i> 'th array element is the index in array <i>pin_GID</i> where the pins for the <i>i</i> 'th row (if <i>format</i> is <code>ZOLTAN_COMPRESSED_EDGE</code> ) or <i>i</i> 'th column (if <i>format</i> is <code>ZOLTAN_COMPRESSED_VERTEX</code> ) begins. Array indices begin at zero.

*pin\_GID* Upon return, a list of *num\_pins* global IDs. This is the list of the pins (or matrix non-zeros) contained in the rows or columns listed in *vtxedge\_GID*.

*ierr* Error code to be set by function.

## Example

	vertex				
hyperedge	10	20	30	40	50
1	0	0	1	1	0
2	0	1	1	0	0
3	1	0	0	0	1

Compressed hyperedge storage:

vtxedge\_GID = { 1, 2, 3 }  
vtxedge\_ptr = { 0, 2, 4 }  
pin\_GID = { 30, 40, 20, 30, 10, 50 }

Compressed vertex storage:

vtxedge\_GID = { 10, 20, 30, 40, 50 }  
vtxedge\_ptr = { 0, 1, 2, 4, 5 }  
pin\_GID = { 3, 2, 1, 2, 1, 3 }

C and C++: `typedef void ZOLTAN_HG_SIZE_EDGE_WTS_FN (void *data, int *num_edges, int *ierr);`

FORTRAN: `SUBROUTINE Get_HG_Size_Edge_Wts(data, num_edges, ierr)`  
`<type-data>, INTENT(IN) :: data`  
`INTEGER(Zoltan_INT), INTENT(OUT) :: num_edges`  
`INTEGER(Zoltan_INT), INTENT(OUT) :: ierr`

`<type-data>` can be any of `INTEGER(Zoltan_INT)`, `DIMENSION(*)` or `REAL(Zoltan_FLOAT)`, `DIMENSION(*)` or `REAL(Zoltan_DOUBLE)`, `DIMENSION(*)` or `TYPE(Zoltan_User_Data_x)` where *x* is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

**A `ZOLTAN_HG_SIZE_EDGE_WTS_FN` returns the number of hyperedges for which a process will supply edge weights. The number of weights per hyperedge was supplied by the application with the [EDGE\\_WEIGHT\\_DIM](#) parameter. The actual edge weights will be supplied with a [ZOLTAN\\_HG\\_EDGE\\_WTS\\_FN\\_TYPE](#) function.**

**This query function is not required. If no hyperedge weights are supplied, Zoltan will assume every hyperedge has weight 1.0.**

**Function Type:** `ZOLTAN_HG_SIZE_EDGE_WTS_FN_TYPE`

**Arguments:**

*data* Pointer to user-defined data.

*num\_edges* Upon return, the number of hyperedges for which edge weights will be supplied.

*ierr* Error code to be set by function.

C and C++: `typedef void ZOLTAN_HG_EDGE_WTS_FN (void *data, int num_gid_entries, int num_lid_entries, int num_edges, int edge_weight_dim, ZOLTAN\_ID\_PTR edge_GID, ZOLTAN\_ID\_PTR edge_LID, float *edge_weight, int *ierr);`

FORTTRAN: `SUBROUTINE Get_HG_Edge_Wts(data, num_gid_entries, num_lid_entries, num_edges, edge_weight_dim, edge_GID, edge_LID, edge_weight, ierr)  
<type-data>, INTENT(IN) :: data  
INTEGER(Zoltan_INT), INTENT(IN) :: num_gid_entries, num_lid_entries, num_edges, edge_weight_dim  
INTEGER(Zoltan_INT), INTENT(OUT), DIMENSION(*) :: edge_GID  
INTEGER(Zoltan_INT), INTENT(OUT), DIMENSION(*) :: edge_LID  
REAL(Zoltan_FLOAT), INTENT(OUT), DIMENSION(*) :: edge_weight  
INTEGER(Zoltan_INT), INTENT(OUT) :: ierr`

<type-data> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE), DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where x is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

---

**A `ZOLTAN_HG_EDGE_WTS_FN` returns edges weights for a set of hypergraph edges. The number of weights supplied for each hyperedge should equal the value of the [EDGE\\_WEIGHT\\_DIM](#) parameter. In the case of a hypergraph which is distributed across multiple processes, if more than one process supplies edge weights for the same hyperedge, the different edge weights will be resolved according to the value of the [PHG\\_EDGE\\_WEIGHT\\_OPERATION](#) parameter.**

**This query function is not required. If no hyperedge weights are supplied, Zoltan will assume every hyperedge has weight 1.0.**

**Function Type:** `ZOLTAN_HG_SIZE_EDGE_WTS_FN_TYPE`

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>num_edges</i>	The number of hyperedges for which edge weights should be supplied in the <i>edge_weight</i> array.
<i>edge_weight_dim</i>	The number of weights which should be supplied for each hyperedge. This is also the value of the <a href="#">EDGE_WEIGHT_DIM</a> parameter.
<i>edge_GID</i>	Upon return, this array should contain the global IDs of the <i>num_edges</i> hyperedges for which the application is supplying edge weights.
<i>edge_LID</i>	Upon return, this array can optionally contain the local IDs of the <i>num_edges</i> hyperedges for which the application is supplying edge weights.
<i>edge_weight</i>	Upon return, this array should contain the weights for each edge listed in the <i>edge_GID</i> . If <i>edge_weight_dim</i> is greater than one, all weights for one hyperedge are listed before the weights for the next hyperedge are listed.
<i>ierr</i>	Error code to be set by function.

---

## Tree-based Functions

---

C and C++:    `typedef int ZOLTAN_NUM_COARSE_OBJ_FN (void *data, int *ierr);`  
 FORTRAN:    `FUNCTION Get_Num_Coarse_Obj(data, ierr)  
 INTEGER(Zoltan_INT) :: Get_Num_Coarse_Obj  
 <type-data>, INTENT(IN) :: data  
 INTEGER(Zoltan_INT), INTENT(OUT) :: ierr`

<*type-data*> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE), DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_*x*) where *x* is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

---

**A ZOLTAN\_NUM\_COARSE\_OBJ\_FN query function returns the number of objects (elements) in the initial coarse grid.**

**Function Type:**    **ZOLTAN\_NUM\_COARSE\_OBJ\_FN\_TYPE**

**Arguments:**

*data*                      Pointer to user-defined data.  
*ierr*                      Error code to be set by function.

**Returned Value:**

int                      The number of objects in the coarse grid.

---

C and C++:    `typedef void ZOLTAN_COARSE_OBJ_LIST_FN (void *data, int num_gid_entries,  
 int num_lid_entries, ZOLTAN\_ID\_PTR global_ids, ZOLTAN\_ID\_PTR local_ids,  
 int *assigned, int *num_vert, ZOLTAN\_ID\_PTR vertices, int *in_order,  
ZOLTAN\_ID\_PTR in_vertex, ZOLTAN\_ID\_PTR out_vertex, int *ierr);`  
 FORTRAN:    `SUBROUTINE Get_Coarse_Obj_List(data, num_gid_entries, num_lid_entries,  
global_ids, local_ids, assigned, num_vert, vertices, in_order, in_vertex, out_vertex,  
ierr)  
 <type-data>, INTENT(IN) :: data  
 INTEGER(Zoltan_INT), INTENT(IN) :: num_gid_entries, num_lid_entries  
 INTEGER(Zoltan_INT), INTENT(OUT), DIMENSION(*) :: global_ids  
 INTEGER(Zoltan_INT), INTENT(OUT), DIMENSION(*) :: local_ids  
 INTEGER(Zoltan_INT), INTENT(OUT), DIMENSION(*) :: assigned, num_vert,  
vertices, in_vertex, out_vertex  
 INTEGER(Zoltan_INT), INTENT(OUT) :: in_order, ierr`

<*type-data*> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE), DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_*x*) where *x* is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

---

**A ZOLTAN\_COARSE\_OBJ\_LIST\_FN query function returns lists of global IDs, local IDs, vertices, and order information for all objects (elements) of the initial coarse grid. The vertices are designated by a global ID such that if two elements share a vertex then the same ID**

designates that vertex in both elements and on all processors. The user may choose to provide the order in which the elements should be traversed or have Zoltan determine the order. If the user provides the order, then entry and exit vertices for a path through the elements may also be provided. The arrays for the returned values are allocated by Zoltan; their size is determined by a call to a [ZOLTAN\\_NUM\\_COARSE\\_OBJ\\_FN](#) query function.

**Function Type:** ZOLTAN\_COARSE\_OBJ\_LIST\_FN\_TYPE

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>global_ids</i>	Upon return, an array of global IDs of all objects in the coarse grid.
<i>local_ids</i>	Upon return, an array of local IDs of all objects in the coarse grid.
<i>assigned</i>	Upon return, an array of integers indicating whether or not each object is currently assigned to this processor. A value of 1 indicates it is assigned to this processor; a value of 0 indicates it is assigned to some other processor. For elements that have been refined, it is ignored unless weights are assigned to interior nodes of the tree.
<i>num_vert</i>	Upon return, an array containing the number of vertices for each object.
<i>vertices</i>	Upon return, an array of global IDs of the vertices of each object. If the number of vertices for objects 0 through $i-1$ is $N$ , then the vertices for object $i$ are in $vertices[N * num\_gid\_entries : (N + num\_vert[i]) * num\_gid\_entries]$
<i>in_order</i>	Upon return, 1 if the user is providing the objects in the order in which they should be traversed, or 0 if Zoltan should determine the order.
<i>in_vertex</i>	Upon return, an array of global IDs of the vertices through which to enter each element in the user provided traversal. It is required only if the user is providing the order for the coarse grid objects (i.e., $in\_order == 1$ ) and allowing Zoltan to select the order of the children in at least one invocation of <a href="#">ZOLTAN_CHILD_LIST_FN</a> .
<i>out_vertex</i>	Upon return, an array of global IDs of the vertex through which to exit each element in the user provided traversal. The same provisions hold as for <i>in_vertex</i> .
<i>ierr</i>	Error code to be set by function.

---

C and C++: `typedef int ZOLTAN_FIRST_COARSE_OBJ_FN (void *data, int num_gid_entries, int num_lid_entries, ZOLTAN\_ID\_PTR global_id, ZOLTAN\_ID\_PTR local_id, int *assigned, int *num_vert, ZOLTAN\_ID\_PTR vertices, int *in_order, ZOLTAN\_ID\_PTR in_vertex, ZOLTAN\_ID\_PTR out_vertex, int *ierr);`

FORTTRAN: `FUNCTION Get_First_Coarse_Obj(data, num_gid_entries, num_lid_entries, global_id, local_id, assigned, num_vert, vertices, in_order, in_vertex, out_vertex, ierr) INTEGER(Zoltan_INT) :: Get_First_Coarse_Obj`  
`<type-data>, INTENT(IN) :: data`  
`INTEGER(Zoltan_INT), INTENT(IN) :: num_gid_entries, num_lid_entries`  
`INTEGER(Zoltan_INT), INTENT(OUT), DIMENSION(*) :: global_id`  
`INTEGER(Zoltan_INT), INTENT(OUT), DIMENSION(*) :: local_id`  
`INTEGER(Zoltan_INT), INTENT(OUT) :: assigned, num_vert, in_order, ierr`  
`INTEGER(Zoltan_INT), INTENT(OUT), DIMENSION(*) :: vertices, in_vertex, out_vertex`

`<type-data>` can be any of `INTEGER(Zoltan_INT)`, `DIMENSION(*)` or `REAL(Zoltan_FLOAT)`, `DIMENSION(*)` or `REAL(Zoltan_DOUBLE)`, `DIMENSION(*)` or `TYPE(Zoltan_User_Data_x)` where  $x$  is 1, 2, 3 or 4. See the section



---

**A ZOLTAN\_FIRST\_COARSE\_OBJ\_FN query function initializes an iteration over the objects of the initial coarse grid. It returns the global ID, local ID, vertices, and order information for the first object (element) of the initial coarse grid. Subsequent calls to a [ZOLTAN\\_NEXT\\_COARSE\\_OBJ\\_FN](#) iterate over and return other objects from the coarse grid. The vertices are designated by a global ID such that if two elements share a vertex then the same ID designates that vertex in both elements and on all processors. The user may choose to provide the order in which the elements should be traversed, or have Zoltan determine the order. If the user provides the order, then entry and exit vertices for a path through the elements may also be provided.**

**Function Type:** ZOLTAN\_FIRST\_COARSE\_OBJ\_FN\_TYPE

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>global_ids</i>	Upon return, the global ID of the first object in the coarse grid.
<i>local_ids</i>	Upon return, the local ID of the first object in the coarse grid.
<i>assigned</i>	Upon return, an integer indicating whether or not this object is currently assigned to this processor. A value of 1 indicates it is assigned to this processor; a value of 0 indicates it is assigned to some other processor. For elements that have been refined, it is ignored unless weights are assigned to interior nodes of the tree.
<i>num_vert</i>	Upon return, the number of vertices for this object.
<i>vertices</i>	Upon return, an array of global IDs of the vertices of this object.
<i>in_order</i>	Upon return, 1 if the user is providing the objects in the order in which they should be traversed, or 0 if Zoltan should determine the order.
<i>in_vertex</i>	Upon return, the vertex through which to enter this element in the user provided traversal. It is required only if the user is providing the order for the coarse grid objects (i.e., <i>in_order</i> ==1) and allowing Zoltan to select the order of the children in at least one invocation of <a href="#">ZOLTAN_CHILD_LIST_FN</a> .
<i>out_vertex</i>	Upon return, the vertex through which to exit this element in the user provided traversal. The same provisions hold as for <i>in_vertex</i> .
<i>ierr</i>	Error code to be set by function.

**Returned Value:**

1	If <i>global_id</i> and <i>local_id</i> contain valid IDs of the first object in the coarse grid.
0	If no coarse grid is available.

---

C and C++: 

```
typedef int ZOLTAN_NEXT_COARSE_OBJ_FN (void *data, int num_gid_entries,
int num_lid_entries, ZOLTAN\_ID\_PTR global_id, ZOLTAN\_ID\_PTR local_id,
ZOLTAN\_ID\_PTR next_global_id, ZOLTAN\_ID\_PTR next_local_id, int *assigned,
int *num_vert, ZOLTAN\_ID\_PTR vertices, ZOLTAN\_ID\_PTR in_vertex,
ZOLTAN\_ID\_PTR out_vertex, int *ierr);
```



**FORTRAN:**    **FUNCTION *Get\_Next\_Coarse\_Obj***(*data*, *num\_gid\_entries*, *num\_lid\_entries*, *global\_id*, *local\_id*, *next\_global\_id*, *next\_local\_id*, *assigned*, *num\_vert*, *vertices*, *in\_vertex*, *out\_vertex*, *ierr*)  
 INTEGER(Zoltan\_INT) :: *Get\_Next\_Coarse\_Obj*  
 <type-data>, INTENT(IN) :: *data*  
 INTEGER(Zoltan\_INT), INTENT(IN) :: *num\_gid\_entries*, *num\_lid\_entries*  
 INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *global\_id*  
 INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *local\_id*  
 INTEGER(Zoltan\_INT), INTENT(OUT), DIMENSION(\*) :: *next\_global\_id*  
 INTEGER(Zoltan\_INT), INTENT(OUT), DIMENSION(\*) :: *next\_local\_id*  
 INTEGER(Zoltan\_INT), INTENT(OUT) :: *assigned*, *num\_vertex*, *ierr*  
 INTEGER(Zoltan\_INT), INTENT(OUT), DIMENSION(\*) :: *vertices*, *in\_vertex*, *out\_vertex*

<type-data> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE), DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

---

**A ZOLTAN\_NEXT\_COARSE\_OBJ\_FN query function is an iterator function that returns the next object in the initial coarse grid. The first object of the iteration is provided by a [ZOLTAN\\_FIRST\\_COARSE\\_OBJ\\_FN](#) query function.**

---

**Function Type:**    **ZOLTAN\_NEXT\_COARSE\_OBJ\_FN\_TYPE**

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>global_id</i>	The global ID of the previous object in the coarse grid.
<i>local_id</i>	The local ID of the previous object in the coarse grid.
<i>next_global_id</i>	Upon return, the global ID of the next object in the coarse grid.
<i>next_local_id</i>	Upon return, the local ID of the next object in the coarse grid.
<i>assigned</i>	Upon return, an integer indicating whether or not this object is currently assigned to this processor. A value of 1 indicates it is assigned to this processor; a value of 0 indicates it is assigned to some other processor. For elements that have been refined, it is ignored unless weights are assigned to interior nodes of the tree.
<i>num_vert</i>	Upon return, the number of vertices for this object.
<i>vertices</i>	Upon return, an array of global IDs of the vertices of this object.
<i>in_vertex</i>	Upon return, the vertex through which to enter this element in the user provided traversal. It is required only if the user is providing the order for the coarse grid objects (i.e., <i>in_order</i> ==1) and allowing Zoltan to select the order of the children in at least one invocation of <a href="#">ZOLTAN_CHILD_LIST_FN</a> .
<i>out_vertex</i>	Upon return, the vertex through which to exit this element in the user provided traversal. The same provisions hold as for <i>in_vertex</i> .
<i>ierr</i>	Error code to be set by function.

**Returned Value:**

1	If <i>global_id</i> and <i>local_id</i> contain valid IDs of the next object in the coarse grid.
0	If no more objects are available.

---

**C and C++:**    typedef int **ZOLTAN\_NUM\_CHILD\_FN** (void \**data*, int *num\_gid\_entries*, int *num\_lid\_entries*, [ZOLTAN\\_ID\\_PTR](#) *global\_id*, [ZOLTAN\\_ID\\_PTR](#) *local\_id*, int \**ierr*);

**FORTRAN:**    **FUNCTION** *Get\_Num\_Child*(*data*, *num\_gid\_entries*, *num\_lid\_entries*, *global\_id*,  
*local\_id*, *ierr*)  
**INTEGER**(Zoltan\_INT) :: *Get\_Num\_Child*  
<*type-data*>, **INTENT**(IN) :: *data*  
**INTEGER**(Zoltan\_INT), **INTENT**(IN) :: *num\_gid\_entries*, *num\_lid\_entries*  
**INTEGER**(Zoltan\_INT), **INTENT**(IN), **DIMENSION**(\*) :: *global\_id*  
**INTEGER**(Zoltan\_INT), **INTENT**(IN), **DIMENSION**(\*) :: *local\_id*  
**INTEGER**(Zoltan\_INT), **INTENT**(OUT) :: *ierr*

<*type-data*> can be any of **INTEGER**(Zoltan\_INT), **DIMENSION**(\*) or **REAL**(Zoltan\_FLOAT), **DIMENSION**(\*) or **REAL**(Zoltan\_DOUBLE), **DIMENSION**(\*) or **TYPE**(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

---

**A ZOLTAN\_NUM\_CHILD\_FN query function returns the number of children of the element with the given global and local IDs. If the element has not been refined, the number of children is 0.**

**Function Type:**    **ZOLTAN\_NUM\_CHILD\_FN\_TYPE**

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>global_id</i>	The global ID of the object for which the number of children is requested.
<i>local_id</i>	The local ID of the object for which the number of children is requested.
<i>ierr</i>	Error code to be set by function.

**Returned Value:**

<i>int</i>	The number of children.
------------	-------------------------

---

**C and C++:**    **typedef void** **ZOLTAN\_CHILD\_LIST\_FN** (**void** \**data*, **int** *num\_gid\_entries*,  
**int** *num\_lid\_entries*, [ZOLTAN\\_ID\\_PTR](#) *parent\_gid*, [ZOLTAN\\_ID\\_PTR](#) *parent\_lid*,  
[ZOLTAN\\_ID\\_PTR](#) *child\_gids*, [ZOLTAN\\_ID\\_PTR](#) *child\_lids*, **int** \**assigned*,  
**int** \**num\_vert*, [ZOLTAN\\_ID\\_PTR](#) *vertices*, **ZOLTAN\_REF\_TYPE** \**ref\_type*,  
[ZOLTAN\\_ID\\_PTR](#) *in\_vertex*, [ZOLTAN\\_ID\\_PTR](#) *out\_vertex*, **int** \**ierr*);

**FORTRAN:**    **SUBROUTINE** *Get\_Child\_List*(*data*, *num\_gid\_entries*, *num\_lid\_entries*, *parent\_gid*,  
*parent\_lid*, *child\_gids*, *child\_lids*, *assigned*, *num\_vert*, *vertices*, *ref\_type*, *in\_vertex*,  
*out\_vertex*, *ierr*)  
<*type-data*>, **INTENT**(IN) :: *data*  
**INTEGER**(Zoltan\_INT), **INTENT**(IN) :: *num\_gid\_entries*, *num\_lid\_entries*  
**INTEGER**(Zoltan\_INT), **INTENT**(IN), **DIMENSION**(\*) :: *parent\_gid*  
**INTEGER**(Zoltan\_INT), **INTENT**(IN), **DIMENSION**(\*) :: *parent\_lid*  
**INTEGER**(Zoltan\_INT), **INTENT**(OUT), **DIMENSION**(\*) :: *child\_gids*  
**INTEGER**(Zoltan\_INT), **INTENT**(OUT), **DIMENSION**(\*) :: *child\_lids*  
**INTEGER**(Zoltan\_INT), **INTENT**(OUT), **DIMENSION**(\*) :: *assigned*, *num\_vert*,  
*vertices*, *in\_vertex*, *out\_vertex*  
**INTEGER**(Zoltan\_INT), **INTENT**(OUT) :: *ref\_type*, *ierr*

<*type-data*> can be any of **INTEGER**(Zoltan\_INT), **DIMENSION**(\*) or **REAL**(Zoltan\_FLOAT), **DIMENSION**(\*) or **REAL**(Zoltan\_DOUBLE), **DIMENSION**(\*) or **TYPE**(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

---

**A ZOLTAN\_CHILD\_LIST\_FN query function returns lists of global IDs,**

**local IDs, vertices, and order information for all children of a refined element. The vertices are designated by a global ID such that if two elements share a vertex then the same ID designates that vertex in both elements and on all processors. The user may choose to provide the order in which the children should be traversed, or have Zoltan determine the order based on the type of element refinement used to create the children. If the user provides the order, then entry and exit vertices for a path through the elements may also be provided. The arrays for the returned values are allocated by Zoltan; their size is determined by a call to a [ZOLTAN\\_NUM\\_CHILD\\_FN](#) query function.**

**Function Type:** ZOLTAN\_CHILD\_LIST\_FN\_TYPE

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>parent_gid</i>	The global ID of the object whose children are requested.
<i>parent_lid</i>	The local ID of the object whose children are requested.
<i>child_gids</i>	Upon return, an array of global IDs of all children of this object.
<i>child_lids</i>	Upon return, an array of local IDs of all children of this object.
<i>assigned</i>	Upon return, an array of integers indicating whether or not each child is currently assigned to this processor. A value of 1 indicates it is assigned to this processor; a value of 0 indicates it is assigned to some other processor. For children that have been further refined, it is ignored unless weights are assigned to interior nodes of the tree.
<i>num_vert</i>	Upon return, an array containing the number of vertices for each object.
<i>vertices</i>	Upon return, an array of global IDs of the vertices of each object. If the number of vertices for objects 0 through $i-1$ is $N$ , then the vertices for object $i$ are in $vertices[N * num\_gid\_entries : (N + num\_vert[i]) * num\_gid\_entries]$
<i>ref_type</i>	Upon return, a value indicating what type of refinement was used to create the children. This determines how the children will be ordered. The values currently supported are: <a href="#">ZOLTAN_TRI_BISECT</a> Bisection of triangles. <a href="#">ZOLTAN_QUAD_QUAD</a> Quadrisection of quadrilaterals. <a href="#">ZOLTAN_HEX3D_OCT</a> Octasection of hexahedra. <a href="#">ZOLTAN_OTHER_REF</a> All other forms of refinement. <a href="#">ZOLTAN_IN_ORDER</a> Traverse the children in the order in which they are provided.
<i>in_vertex</i>	Upon return, an array of global IDs of the vertex through which to enter each element in the user provided traversal. It is required only if the user is providing the order for the children of this element (i.e., $ref\_type == ZOLTAN\_IN\_ORDER$ ) but does not provide the order for the children of at least one of those children.
<i>out_vertex</i>	Upon return, an array of global IDs of the vertex through which to exit each element in the user provided traversal. The same provisions hold as for <i>in_vertex</i> .
<i>ierr</i>	Error code to be set by function.

---

C and C++: `typedef void ZOLTAN\_CHILD\_WEIGHT\_FN (void *data, int num_gid_entries, int num_lid_entries, ZOLTAN\_ID\_PTR global_id, ZOLTAN\_ID\_PTR local_id, int wgt_dim, float *obj_wgt, int *ierr);`

**FORTTRAN:** SUBROUTINE *Get\_Child\_Weight*(*data*, *num\_gid\_entries*, *num\_lid\_entries*, *global\_id*, *local\_id*, *wgt\_dim*, *obj\_wgt*, *ierr*)  
 <type-data>, INTENT(IN) :: *data*  
 INTEGER(Zoltan\_INT), INTENT(IN) :: *num\_gid\_entries*, *num\_lid\_entries*  
 INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *global\_id*  
 INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *local\_id*  
 INTEGER(Zoltan\_INT), INTENT(IN) :: *wgt\_dim*  
 REAL(Zoltan\_FLOAT), INTENT(OUT), DIMENSION(\*) :: *obj\_wgt*  
 INTEGER(Zoltan\_INT), INTENT(OUT) :: *ierr*

<type-data> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE), DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

---

**A ZOLTAN\_CHILD\_WEIGHT\_FN query function returns the weight of an object. Interior nodes of the refinement tree as well as the leaves are allowed to have weights.**

**Function Type:** ZOLTAN\_CHILD\_WEIGHT\_FN\_TYPE

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>global_id</i>	The global ID of the object whose weight is requested.
<i>local_id</i>	The local ID of the object whose weight is requested.
<i>wgt_dim</i>	The number of weights associated with an object (typically 1), or 0 if weights are not requested. This value is set through the parameter <a href="#">OBJ_WEIGHT_DIM</a> .
<i>obj_wgt</i>	Upon return, an array containing the object's weights. If <i>wgt_dim</i> =0, the return value of <i>obj_wgt</i> s is undefined and may be NULL.
<i>ierr</i>	Error code to be set by function.

## Border Object Functions (currently not used)

---

**C:** typedef int ZOLTAN\_NUM\_BORDER\_OBJ\_FN (void \**data*, int *nbor\_proc*, int \**ierr*);

**FORTTRAN:** FUNCTION *Get\_Num\_Border\_Obj*(*data*, *nbor\_proc*, *ierr*)  
 INTEGER(Zoltan\_INT) :: *Get\_Num\_Border\_Obj*  
 <type-data>, INTENT(IN) :: *data*  
 INTEGER(Zoltan\_INT), INTENT(IN) :: *nbor\_proc*  
 INTEGER(Zoltan\_INT), INTENT(OUT) :: *ierr*

<type-data> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE), DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

---

**A ZOLTAN\_NUM\_BORDER\_OBJ\_FN query function returns the number of objects sharing a processor subdomain border (in the communication graph of the application) with a given processor.**

**Function Type:** ZOLTAN\_NUM\_BORDER\_OBJ\_FN\_TYPE

**Arguments:**

*data* Pointer to user-defined data.

*nbor\_proc* The processor ID of the processor for which the number of border objects should be returned.

*ierr* Error code to be set by function.

**Returned Value:**

int The number of objects sharing a processor subdomain border with processor *nbor\_proc*.

C: typedef void **ZOLTAN\_BORDER\_OBJ\_LIST\_FN** (void \**data*, int *num\_gid\_entries*, int *num\_lid\_entries*, int *nbor\_proc*, [ZOLTAN\\_ID\\_PTR](#) *global\_ids*, [ZOLTAN\\_ID\\_PTR](#) *local\_ids*, int *wgt\_dim*, float \**obj\_wgts*, int \**ierr*);

FORTRAN: SUBROUTINE *Get\_Border\_Obj\_List*(*data*, *num\_gid\_entries*, *num\_lid\_entries*, *nbor\_proc*, *global\_ids*, *local\_ids*, *wgt\_dim*, *obj\_wgts*, *ierr*)  
 <type-data>, INTENT(IN) :: *data*  
 INTEGER(Zoltan\_INT), INTENT(IN) :: *nbor\_proc*  
 INTEGER(Zoltan\_INT), INTENT(IN) :: *num\_gid\_entries*, *num\_lid\_entries*  
 INTEGER(Zoltan\_INT), INTENT(OUT), DIMENSION(\*) :: *global\_ids*  
 INTEGER(Zoltan\_INT), INTENT(OUT), DIMENSION(\*) :: *local\_ids*  
 INTEGER(Zoltan\_INT), INTENT(IN) :: *wgt\_dim*  
 REAL(Zoltan\_FLOAT), INTENT(OUT), DIMENSION(\*) :: *obj\_wgts*  
 INTEGER(Zoltan\_INT), INTENT(OUT) :: *ierr*

<type-data> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE), DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

**A ZOLTAN\_BORDER\_OBJ\_LIST\_FN query function fills two arrays with information about the objects currently assigned to the processor that share a processor subdomain border (in the communication graph of the application) with a given processor. Both arrays are allocated (and subsequently freed) by Zoltan; their size is determined by a call to a [ZOLTAN\\_NUM\\_BORDER\\_OBJ\\_FN](#) query function to get the array size. For certain Zoltan algorithms, either a ZOLTAN\_BORDER\_OBJ\_LIST\_FN query function or a [ZOLTAN\\_FIRST\\_BORDER\\_OBJ\\_FN/ZOLTAN\\_NEXT\\_BORDER\\_OBJ\\_FN](#) query-function pair must be registered; however, both query options need not be provided.**

**Function Type:** ZOLTAN\_BORDER\_OBJ\_LIST\_FN\_TYPE

**Arguments:**

*data* Pointer to user-defined data.

*num\_gid\_entries* The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter [NUM\\_GID\\_ENTRIES](#).

*num\_lid\_entries* The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter [NUM\\_LID\\_ENTRIES](#).

*nbor\_proc* The processor ID of the processor for which border objects should be returned.



<i>global_ids</i>	Upon return, an array of unique global IDs for all objects assigned to the processor that share a subdomain border with <i>nbor_proc</i> .
<i>local_ids</i>	Upon return, an array of local IDs, the meaning of which can be determined by the application, for all objects assigned to the processor that share a subdomain border with <i>nbor_proc</i> .
<i>wgt_dim</i>	The number of weights associated with an object (typically 1), or 0 if weights are not requested. This value is set through the parameter <a href="#">OBJ_WEIGHT_DIM</a> .
<i>obj_wgts</i>	Upon return, an array of object weights. Weights for object <i>i</i> are stored in <i>obj_wgts[(i-1)*wgt_dim:i*wgt_dim-1]</i> . If <i>wgt_dim=0</i> , <i>obj_wgts</i> is undefined and may be NULL.
<i>ierr</i>	Error code to be set by function.

---

C:           typedef int **ZOLTAN\_FIRST\_BORDER\_OBJ\_FN** (void \*data, int num\_gid\_entries, int num\_lid\_entries, int nbor\_proc, [ZOLTAN\\_ID\\_PTR](#) first\_global\_id, [ZOLTAN\\_ID\\_PTR](#) first\_local\_id, int wgt\_dim, float \*first\_obj\_wgt, int \*ierr);

FORTRAN:   FUNCTION *Get\_First\_Border\_Obj*(data, num\_gid\_entries, num\_lid\_entries, nbor\_proc, first\_global\_id, first\_local\_id, wgt\_dim, first\_obj\_wgt, ierr)  
 INTEGER(Zoltan\_INT) :: Get\_First\_Border\_Obj  
 <type-data>, INTENT(IN) :: data  
 INTEGER(Zoltan\_INT), INTENT(IN) :: num\_gid\_entries, num\_lid\_entries  
 INTEGER(Zoltan\_INT), INTENT(IN) :: nbor\_proc  
 INTEGER(Zoltan\_INT), INTENT(OUT), DIMENSION(\*) :: first\_global\_id  
 INTEGER(Zoltan\_INT), INTENT(OUT), DIMENSION(\*) :: first\_local\_id  
 INTEGER(Zoltan\_INT), INTENT(IN) :: wgt\_dim  
 REAL(Zoltan\_FLOAT), INTENT(OUT), DIMENSION(\*) :: first\_obj\_wgt  
 INTEGER(Zoltan\_INT), INTENT(OUT) :: ierr

<type-data> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE), DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

---

**A ZOLTAN\_FIRST\_BORDER\_OBJ\_FN query function initializes an iteration over objects assigned to the processor that share a processor subdomain border with a given processor. It returns the global and local IDs of the first object on the processor along the specified subdomain border. Subsequent calls to a [ZOLTAN\\_NEXT\\_BORDER\\_OBJ\\_FN](#) query function iterate over and return other objects along the requested subdomain border. This query-function pair frees the application from having to build an array of objects (as in [ZOLTAN\\_BORDER\\_OBJ\\_LIST\\_FN](#)) and allows Zoltan to obtain only as much information about objects as it needs. For some algorithms, either a [ZOLTAN\\_BORDER\\_OBJ\\_LIST\\_FN](#) query function or a [ZOLTAN\\_FIRST\\_BORDER\\_OBJ\\_FN/ZOLTAN\\_NEXT\\_BORDER\\_OBJ\\_FN](#) query-function pair must be registered; however, both query options need not be provided.**

**Function Type:** **ZOLTAN\_FIRST\_BORDER\_OBJ\_FN\_TYPE****Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>nbor_proc</i>	The processor ID of the processor for which border objects should be returned.
<i>first_global_id</i>	The returned value of the global ID for the first object; the value is ignored if there are no objects along the border.
<i>first_local_id</i>	The returned value of the local ID for the first object; the value is ignored if there are no objects along the border.
<i>wgt_dim</i>	The number of weights associated with an object (typically 1), or 0 if weights are not requested. This value is set through the parameter <a href="#">OBJ_WEIGHT_DIM</a> .
<i>first_obj_wgt</i>	Upon return, the first object's weights; an array of size <i>wgt_dim</i> . Undefined if <i>wgt_dim</i> =0.
<i>ierr</i>	Error code to be set by function.

**Returned Value:**

1	If <i>first_global_id</i> and <i>first_local_id</i> contain valid IDs of the first object along the processor border.
0	If no objects are available along this processor border.

C:       typedef int **ZOLTAN\_NEXT\_BORDER\_OBJ\_FN** (void \**data*, int *num\_gid\_entries*,  
int *num\_lid\_entries*, [ZOLTAN\\_ID\\_PTR](#) *global\_id*, [ZOLTAN\\_ID\\_PTR](#) *local\_id*,  
int *nbor\_proc*, [ZOLTAN\\_ID\\_PTR](#) *next\_global\_id*, [ZOLTAN\\_ID\\_PTR](#) *next\_local\_id*,  
int *wgt\_dim*, float \**next\_obj\_wgt*, int \**ierr*);

FORTTRAN:   FUNCTION *Get\_Next\_Border\_Obj*(*data*, *num\_gid\_entries*, *num\_lid\_entries*, *global\_id*,  
*local\_id*, *nbor\_proc*, *next\_global\_id*, *next\_local\_id*, *wgt\_dim*, *next\_obj\_wgt*, *ierr*)  
INTEGER(Zoltan\_INT) :: *Get\_Next\_Border\_Obj*  
<*type-data*>, INTENT(IN) :: *data*  
INTEGER(Zoltan\_INT), INTENT(IN) :: *num\_gid\_entries*, *num\_lid\_entries*  
INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *global\_id*  
INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *local\_id*  
INTEGER(Zoltan\_INT), INTENT(IN) :: *nbor\_proc*  
INTEGER(Zoltan\_INT), INTENT(OUT), DIMENSION(\*) :: *next\_global\_id*  
INTEGER(Zoltan\_INT), INTENT(OUT), DIMENSION(\*) :: *next\_local\_id*  
INTEGER(Zoltan\_INT), INTENT(IN) :: *wgt\_dim*  
REAL(Zoltan\_FLOAT), INTENT(OUT), DIMENSION(\*) :: *next\_obj\_wgt*  
INTEGER(Zoltan\_INT), INTENT(OUT) :: *ierr*

<*type-data*> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or  
REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE),  
DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section  
on [Fortran query functions](#) for an explanation.

**A ZOLTAN\_NEXT\_BORDER\_OBJ\_FN query function is an iterator function which, when given an object assigned to the processor and a neighboring processor ID, returns the next object assigned to the processor that shares a subdomain border with the neighboring processor. The first object of the iteration is provided by a [ZOLTAN\\_FIRST\\_BORDER\\_OBJ\\_FN](#) query function. This query-function pair frees the application from having to build an array of objects (as in [ZOLTAN\\_BORDER\\_OBJ\\_LIST\\_FN](#)) and allows Zoltan**



to obtain only as much information about objects as it needs. For some algorithms, either a [ZOLTAN\\_BORDER\\_OBJ\\_LIST\\_FN](#) query function or a [ZOLTAN\\_FIRST\\_BORDER\\_OBJ\\_FN/ZOLTAN\\_NEXT\\_BORDER\\_OBJ\\_FN](#) query-function pair must be registered; however, both query options need not be provided.

<b>Function Type:</b>	<b>ZOLTAN_NEXT_BORDER_OBJ_FN_TYPE</b>
<b>Arguments:</b>	
<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>global_id</i>	The global ID of the previous object.
<i>local_id</i>	The local ID of the previous object.
<i>nbor_proc</i>	The processor ID of the processor for which border objects should be returned.
<i>next_global_id</i>	The returned value of the global ID for the next object; the value is ignored if there are no more objects along the border.
<i>next_local_id</i>	The returned value of the local ID for the next object; the value is ignored if there are no more objects along the border.
<i>wgt_dim</i>	The number of weights associated with an object (typically 1), or 0 if weights are not requested. This value is set through the parameter <a href="#">OBJ_WEIGHT_DIM</a> .
<i>next_obj_wgt</i>	Upon return, the weights for the next object; an array of size <i>wgt_dim</i> . Undefined if <i>wgt_dim</i> =0.
<i>ierr</i>	Error code to be set by function.
<b>Returned Value:</b>	
1	If <i>next_global_id</i> and <i>next_local_id</i> contain valid IDs of the next object along the processor border.
0	If no more objects are available along this processor border.

---

[\[Table of Contents](#) | [Next: Migration Query Functions](#) | [Previous: Application-Registered Query Functions\]](#)

# Migration Query Functions

The following query functions must be registered to use any of the migration tools described in [Migration Functions](#):

[ZOLTAN\\_OBJ\\_SIZE\\_FN](#) or [ZOLTAN\\_OBJ\\_SIZE\\_MULTI\\_FN](#)  
[ZOLTAN\\_PACK\\_OBJ\\_FN](#) or [ZOLTAN\\_PACK\\_OBJ\\_MULTI\\_FN](#)  
[ZOLTAN\\_UNPACK\\_OBJ\\_FN](#) or [ZOLTAN\\_UNPACK\\_OBJ\\_MULTI\\_FN](#)

The "MULTI\_" versions of the packing/unpacking functions take lists of IDs as input and pack/unpack data for all objects in the lists. Only one function of each type must be provided (e.g., either a [ZOLTAN\\_PACK\\_OBJ\\_FN](#) or [ZOLTAN\\_PACK\\_OBJ\\_MULTI\\_FN](#), but not both).

Optional, additional query functions for migration may also be registered; these functions are called at the beginning, middle, and end of migration in [Zoltan\\_Migrate](#).

[ZOLTAN\\_PRE\\_MIGRATE\\_PP\\_FN](#)  
[ZOLTAN\\_MID\\_MIGRATE\\_PP\\_FN](#)  
[ZOLTAN\\_POST\\_MIGRATE\\_PP\\_FN](#)

For [backward compatibility](#) with previous versions of Zoltan, the following functions may be used with [Zoltan\\_Help\\_Migrate](#).

[ZOLTAN\\_PRE\\_MIGRATE\\_FN](#)  
[ZOLTAN\\_MID\\_MIGRATE\\_FN](#)  
[ZOLTAN\\_POST\\_MIGRATE\\_FN](#)

---

```
C and C++:    typedef int ZOLTAN_OBJ_SIZE_FN(
                void *data,
                int num_gid_entries,
                int num_lid_entries,
                ZOLTAN\_ID\_PTR global_id,
                ZOLTAN\_ID\_PTR local_id,
                int *ierr);
```

**FORTTRAN:**     **FUNCTION** *Obj\_Size*(*data*, *num\_gid\_entries*, *num\_lid\_entries*, *global\_id*, *local\_id*,  
*ierr*)  
INTEGER(Zoltan\_INT) :: Obj\_Size  
<*type-data*>, INTENT(IN) :: *data*  
INTEGER(Zoltan\_INT), INTENT(IN) :: *num\_gid\_entries*, *num\_lid\_entries*  
INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *global\_id*, *local\_id*  
INTEGER(Zoltan\_INT), INTENT(OUT) :: *ierr*

<*type-data*> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE), DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

---

A **ZOLTAN\_OBJ\_SIZE\_FN** query function returns the size (in bytes) of the data buffer that is needed to pack all of a single object's data.

**Function Type:**     **ZOLTAN\_OBJ\_SIZE\_FN\_TYPE**

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>global_id</i>	Pointer to the global ID of the object.
<i>local_id</i>	Pointer to the local ID of the object.
<i>ierr</i>	Error code to be set by function.

**Returned Value:**

int	The size (in bytes) of the required data buffer.
-----	--

---

**C and C++:**     typedef void **ZOLTAN\_OBJ\_SIZE\_MULTI\_FN** (  
                    void \**data*,  
                    int *num\_gid\_entries*,  
                    int *num\_lid\_entries*,  
                    int *num\_ids*,  
                    [ZOLTAN\\_ID\\_PTR](#) *global\_ids*,  
                    [ZOLTAN\\_ID\\_PTR](#) *local\_ids*,  
                    int \**sizes*,  
                    int \**ierr*);

**FORTTRAN:**     **SUBROUTINE *Obj\_Size\_Multi***(*data*, *num\_gid\_entries*, *num\_lid\_entries*, *num\_ids*,  
                   *global\_ids*, *local\_ids*, *sizes*, *ierr*)  
                   <*type-data*>, INTENT(IN) :: *data*  
                   INTEGER(Zoltan\_INT), INTENT(IN) :: *num\_gid\_entries*, *num\_lid\_entries*, *num\_ids*  
                   INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *global\_ids*, *local\_ids*  
                   INTEGER(Zoltan\_INT), INTENT(OUT), DIMENSION(\*) :: *sizes*  
                   INTEGER(Zoltan\_INT), INTENT(OUT) :: *ierr*

<*type-data*> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE), DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

---

A **ZOLTAN\_OBJ\_SIZE\_MULTI\_FN** query function is the multiple-ID version of [ZOLTAN\\_OBJ\\_SIZE\\_FN](#). For a list of objects, it returns the per-objects sizes (in bytes) of the data buffers needed to pack object data.

**Function Type:**     **ZOLTAN\_OBJ\_SIZE\_MULTI\_FN\_TYPE**

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>num_ids</i>	The number of objects whose sizes are to be returned.
<i>global_ids</i>	An array of global IDs of the objects. The ID for the <i>i</i> -th object begins in <i>global_ids</i> [ <i>i</i> * <i>num_gid_entries</i> ].
<i>local_ids</i>	An array of local IDs of the objects. The ID for the <i>i</i> -th object begins in <i>local_ids</i> [ <i>i</i> * <i>num_lid_entries</i> ].
<i>sizes</i>	Upon return, array of sizes (in bytes) for each object in the ID lists.
<i>ierr</i>	Error code to be set by function.

**Returned Value:**

int	The size (in bytes) of the required data buffer.
-----	--

---

C and C++:     typedef void **ZOLTAN\_PACK\_OBJ\_FN** (  
                   void *\*data*,  
                   int *num\_gid\_entries*,  
                   int *num\_lid\_entries*,  
                   [ZOLTAN\\_ID\\_PTR](#) *global\_id*,  
                   [ZOLTAN\\_ID\\_PTR](#) *local\_id*,  
                   int *dest*,  
                   int *size*,  
                   char *\*buf*,  
                   int *\*ierr*);

FORTRAN:     SUBROUTINE **Pack\_Obj**(*data*, *num\_gid\_entries*, *num\_lid\_entries*, *global\_id*, *local\_id*,  
                   *dest*, *size*, *buf*, *ierr*)

<type-data>, INTENT(IN) :: data

INTEGER(Zoltan\_INT), INTENT(IN) :: num\_gid\_entries, num\_lid\_entries

INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: global\_id

INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: local\_id

INTEGER(Zoltan\_INT), INTENT(IN) :: dest, size

INTEGER(Zoltan\_INT), INTENT(OUT), DIMENSION(\*) :: buf

INTEGER(Zoltan\_INT), INTENT(OUT) :: ierr

<type-data> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE), DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

---

A **ZOLTAN\_PACK\_OBJ\_FN** query function allows the application to tell Zoltan how to copy all needed data for a given object into a communication buffer. The object's data can then be sent to another processor as part of data migration. It may also perform other operations, such as removing the object from the processor's data structure. This routine is called by [Zoltan\\_Migrate](#) for each object to be sent to another processor.

**Function Type:**     **ZOLTAN\_PACK\_OBJ\_FN\_TYPE**

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>global_id</i>	The global ID of the object for which data should be copied into the communication buffer.
<i>local_id</i>	The local ID of the object for which data should be copied into the communication buffer.
<i>dest</i>	The destination partition (i.e., the partition to which the object is being sent)

<i>size</i>	The size (in bytes) of the communication buffer for the specified object (as returned by the <a href="#">ZOLTAN_OBJ_SIZE_FN</a> query function).
<i>buf</i>	The starting address of the communication buffer into which the object's data should be packed.
<i>ierr</i>	Error code to be set by function.

---

C and C++:     typedef void **ZOLTAN\_PACK\_OBJ\_MULTI\_FN** (  
                   void \**data*,  
                   int *num\_gid\_entries*,  
                   int *num\_lid\_entries*,  
                   int *num\_ids*,  
                   [ZOLTAN\\_ID\\_PTR](#) *global\_ids*,  
                   [ZOLTAN\\_ID\\_PTR](#) *local\_ids*,  
                   int \**dest*,  
                   int \**sizes*,  
                   int \**idx*,  
                   char \**buf*,  
                   int \**ierr*);

FORTTRAN:     SUBROUTINE **Pack\_Obj\_Multi**(*data*, *num\_gid\_entries*, *num\_lid\_entries*, *num\_ids*,  
                   *global\_ids*, *local\_ids*, *dest*, *sizes*, *idx*, *buf*, *ierr*)  
                   <type-data>, INTENT(IN) :: *data*  
                   INTEGER(Zoltan\_INT), INTENT(IN) :: *num\_gid\_entries*, *num\_lid\_entries*, *num\_ids*  
                   INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *global\_ids*  
                   INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *local\_ids*  
                   INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *dest*  
                   INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *sizes*  
                   INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *idx*  
                   INTEGER(Zoltan\_INT), INTENT(OUT), DIMENSION(\*) :: *buf*  
                   INTEGER(Zoltan\_INT), INTENT(OUT) :: *ierr*

                  <type-data> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or  
                   REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE),  
                   DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_*x*) where *x* is 1, 2, 3 or 4. See the section  
                   on [Fortran query functions](#) for an explanation.

---

A **ZOLTAN\_PACK\_OBJ\_MULTI\_FN** query function is the multiple-ID version of a [ZOLTAN\\_PACK\\_OBJ\\_FN](#). It allows the application to tell Zoltan how to copy all needed data for a given list of objects into a communication buffer.

**Function Type:**     **ZOLTAN\_PACK\_OBJ\_FN\_MULTI\_TYPE**

**Arguments:**

*data*             Pointer to user-defined data.

<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>num_ids</i>	The number of objects to be packed.
<i>global_ids</i>	An array of global IDs of the objects. The ID for the <i>i</i> -th object begins in <i>global_ids[i*num_gid_entries]</i> .
<i>local_ids</i>	An array of local IDs of the objects. The ID for the <i>i</i> -th object begins in <i>local_ids[i*num_lid_entries]</i> .
<i>dest</i>	An array of destination partition numbers (i.e., the partitions to which the objects are being sent)
<i>sizes</i>	An array containing the per-object sizes (in bytes) of the communication buffer for each object.
<i>idx</i>	For each object, an index into the <i>buf</i> array giving the starting location of that object's data. Data for the <i>i</i> -th object are stored in <i>buf[idx[i]]</i> , <i>buf[idx[i]+1]</i> , ..., <i>buf[idx[i]+sizes[i]-1]</i> . Because Zoltan adds some tag information to packed data, $idx[i] \neq \sum_{j=0, i-1} (sizes[j])$ .
<i>buf</i>	The address of the communication buffer into which the objects' data should be packed.
<i>ierr</i>	Error code to be set by function.

C and C++:     typedef void **ZOLTAN\_UNPACK\_OBJ\_FN** (  
                   void \**data*,  
                   int *num\_gid\_entries*,  
                   [ZOLTAN\\_ID\\_PTR](#) *global\_id*,  
                   int *size*,  
                   char \**buf*,  
                   int \**ierr*);

FORTRAN:     SUBROUTINE *Unpack\_Obj*(*data*, *num\_gid\_entries*, *global\_id*, *size*, *buf*, *ierr*)  
               <*type-data*>, INTENT(INOUT) :: *data*  
               INTEGER(Zoltan\_INT), INTENT(IN) :: *num\_gid\_entries*  
               INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *global\_id*  
               INTEGER(Zoltan\_INT), INTENT(IN) :: *size*  
               INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *buf*  
               INTEGER(Zoltan\_INT), INTENT(OUT) :: *ierr*

              <*type-data*> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE), DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

A **ZOLTAN\_UNPACK\_OBJ\_FN** query function allows the application to tell Zoltan how to copy all needed data for a given object from a communication buffer into the application's data structure. This operation is needed as the final step of importing objects during data migration. The query function may



also perform other computation, such as building request lists for related data. This routine is called by [Zoltan\\_Migrate](#) for each object to be received by the processor. (Note: a local ID for the object is not included in this function, as the local ID is local to the exporting, not the importing, processor.)

**Function Type:** **ZOLTAN\_UNPACK\_OBJ\_FN\_TYPE**

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>global_id</i>	The global ID of the object whose data has been received in the communication buffer.
<i>size</i>	The size (in bytes) of the object's data in the communication buffer.
<i>buf</i>	The starting address of the communication buffer for this object.
<i>ierr</i>	Error code to be set by function.

C and C++: typedef void **ZOLTAN\_UNPACK\_OBJ\_MULTI\_FN** (  
 void *\*data*,  
 int *num\_gid\_entries*,  
 int *num\_ids*,  
[ZOLTAN\\_ID\\_PTR](#) *global\_ids*,  
 int *\*sizes*,  
 int *\*idx*,  
 char *\*buf*,  
 int *\*ierr*);

FORTRAN: SUBROUTINE *Unpack\_Obj\_Multi*(*data*, *num\_gid\_entries*, *num\_ids*, *global\_ids*, *sizes*,  
*idx*, *buf*, *ierr*)  
 <type-data>, INTENT(INOUT) :: data  
 INTEGER(Zoltan\_INT), INTENT(IN) :: num\_gid\_entries  
 INTEGER(Zoltan\_INT), INTENT(IN) :: num\_ids  
 INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: global\_ids  
 INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: sizes  
 INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: idx  
 INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: buf  
 INTEGER(Zoltan\_INT), INTENT(OUT) :: ierr  
  
 <type-data> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or  
 REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE),  
 DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section  
 on [Fortran query functions](#) for an explanation.

A **ZOLTAN\_UNPACK\_OBJ\_MULTI\_FN** query function is the multiple-ID version of a [ZOLTAN\\_UNPACK\\_OBJ\\_FN](#). It allows the application to tell Zoltan how to copy all needed data for a given list of objects from a communication buffer into the application's data structure.

**Function Type:**     **ZOLTAN\_UNPACK\_OBJ\_MULTI\_FN\_TYPE**

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_ids</i>	The number of objects to be unpacked.
<i>global_ids</i>	An array of global IDs of the objects. The ID for the <i>i</i> -th object begins in <i>global_ids[i*num_gid_entries]</i> .
<i>sizes</i>	An array containing the per-object sizes (in bytes) of the communication buffer for each object.
<i>idx</i>	For each object, an index into the <i>buf</i> array giving the starting location of that object's data. Data for the <i>i</i> -th object are stored in <i>buf[idx[i]]</i> , <i>buf[idx[i]+1]</i> , ..., <i>buf[idx[i]+sizes[i]-1]</i> . Because Zoltan adds some tag information to packed data, $idx[i] \neq \sum_{j=0, i-1} (sizes[j])$ .
<i>buf</i>	The address of the communication buffer from which data is unpacked.
<i>ierr</i>	Error code to be set by function.

---



---

C and C++:     typedef void **ZOLTAN\_PRE\_MIGRATE\_PP\_FN** (  
                   void \**data*,  
                   int *num\_gid\_entries*,  
                   int *num\_lid\_entries*,  
                   int *num\_import*,  
                   [ZOLTAN\\_ID\\_PTR](#) *import\_global\_ids*,  
                   [ZOLTAN\\_ID\\_PTR](#) *import\_local\_ids*,  
                   int \**import\_procs*,  
                   int \**import\_to\_part*,  
                   int *num\_export*,  
                   [ZOLTAN\\_ID\\_PTR](#) *export\_global\_ids*,  
                   [ZOLTAN\\_ID\\_PTR](#) *export\_local\_ids*,  
                   int \**export\_procs*,  
                   int \**export\_to\_part*,  
                   int \**ierr*);

**FORTTRAN:**     **SUBROUTINE *Pre\_Migrate\_PP***(*data*, *num\_gid\_entries*, *num\_lid\_entries*, *num\_import*,  
*import\_global\_ids*, *import\_local\_ids*, *import\_procs*, *import\_to\_part*, *num\_export*,  
*export\_global\_ids*, *export\_local\_ids*, *export\_procs*, *export\_to\_part*, *ierr*)  
<type-data>, INTENT(INOUT) :: *data*  
INTEGER(Zoltan\_INT), INTENT(IN) :: *num\_gid\_entries*, *num\_lid\_entries*  
INTEGER(Zoltan\_INT), INTENT(IN) :: *num\_import*, *num\_export*  
INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *import\_global\_ids*,  
*export\_global\_ids*  
INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *import\_local\_ids*,  
*export\_local\_ids*  
INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *import\_procs*, *export\_procs*  
INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *import\_to\_part*,  
*export\_to\_part*  
INTEGER(Zoltan\_INT), INTENT(OUT) :: *ierr*

<type-data> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or  
REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE),  
DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section  
on [Fortran query functions](#) for an explanation.

---

A **ZOLTAN\_PRE\_MIGRATE\_PP\_FN** query function performs any pre-processing desired by the application. If it is registered, it is called at the beginning of the [Zoltan\\_Migrate](#) routine. The arguments passed to [Zoltan\\_Migrate](#) are made available for use in the pre-processing routine.

**Function Type:**     **ZOLTAN\_PRE\_MIGRATE\_PP\_FN\_TYPE**

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>num_import</i>	The number of objects that will be received by this processor.
<i>import_global_ids</i>	An array of <i>num_import</i> global IDs of objects to be received by this processor. This array may be NULL, as the processor does not necessarily need to know which objects it will receive.
<i>import_local_ids</i>	An array of <i>num_import</i> local IDs of objects to be received by this processor. This array may be NULL, as the processor does not necessarily need to know which objects it will receive.
<i>import_procs</i>	An array of size <i>num_import</i> listing the processor IDs of the source processors. This array may be NULL, as the processor does not necessarily need to know which objects it will receive.
<i>import_to_part</i>	An array of size <i>num_import</i> listing the partitions to which objects will be imported. This array may be NULL, as the processor does not necessarily need to know from which objects it will receive.

<i>num_export</i>	The number of objects that will be sent from this processor to other processors.
<i>export_global_ids</i>	An array of <i>num_export</i> global IDs of objects to be sent from this processor.
<i>export_local_ids</i>	An array of <i>num_export</i> local IDs of objects to be sent from this processor.
<i>export_procs</i>	An array of size <i>num_export</i> listing the processor IDs of the destination processors.
<i>export_to_part</i>	An array of size <i>num_export</i> listing the partitions to which objects will be sent.
<i>ierr</i>	Error code to be set by function.

**Default:**

No pre-processing is done if a **ZOLTAN\_PRE\_MIGRATE\_PP\_FN** is not registered.

C and C++:     typedef void **ZOLTAN\_MID\_MIGRATE\_PP\_FN** (  
                   void \*data,  
                   int num\_gid\_entries,  
                   int num\_lid\_entries,  
                   int num\_import,  
                   [ZOLTAN\\_ID\\_PTR](#) import\_global\_ids,  
                   [ZOLTAN\\_ID\\_PTR](#) import\_local\_ids,  
                   int \*import\_procs,  
                   int \*import\_to\_part,  
                   int num\_export,  
                   [ZOLTAN\\_ID\\_PTR](#) export\_global\_ids,  
                   [ZOLTAN\\_ID\\_PTR](#) export\_local\_ids,  
                   int \*export\_procs,  
                   int \*export\_to\_part,  
                   int \*ierr);

FORTTRAN:     SUBROUTINE *Mid\_Migrate\_PP*(data, num\_gid\_entries, num\_lid\_entries,  
                   num\_import, import\_global\_ids, import\_local\_ids, import\_procs, import\_to\_part,  
                   num\_export, export\_global\_ids, export\_local\_ids, export\_procs, export\_to\_part, ierr)  
                   <type-data>, INTENT(INOUT) :: data  
                   INTEGER(Zoltan\_INT), INTENT(IN) :: num\_gid\_entries, num\_lid\_entries  
                   INTEGER(Zoltan\_INT), INTENT(IN) :: num\_import, num\_export  
                   INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: import\_global\_ids,  
                   export\_global\_ids  
                   INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: import\_local\_ids,  
                   export\_local\_ids  
                   INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: import\_procs, export\_procs  
                   INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: import\_to\_part,  
                   export\_to\_part  
                   INTEGER(Zoltan\_INT), INTENT(OUT) :: ierr

<type-data> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE), DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where x is 1, 2, 3 or 4. See the section

on [Fortran query functions](#) for an explanation.

---

A **ZOLTAN\_MID\_MIGRATE\_PP\_FN** query function performs any processing desired by the application between the packing and unpacking of objects being migrated. If it is registered, it is called after export objects are packed in [Zoltan\\_Migrate](#); imported objects are unpacked after the **ZOLTAN\_MID\_MIGRATE\_PP\_FN** query function is called. The arguments passed to [Zoltan\\_Migrate](#) are made available for use in the processing routine.

**Function Type:** **ZOLTAN\_MID\_MIGRATE\_PP\_FN\_TYPE**

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .
<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>num_import</i>	The number of objects that will be received by this processor.
<i>import_global_ids</i>	An array of <i>num_import</i> global IDs of objects to be received by this processor. This array may be NULL, as the processor does not necessarily need to know which objects it will receive.
<i>import_local_ids</i>	An array of <i>num_import</i> local IDs of objects to be received by this processor. This array may be NULL, as the processor does not necessarily need to know which objects it will receive.
<i>import_procs</i>	An array of size <i>num_import</i> listing the processor IDs of the source processors. This array may be NULL, as the processor does not necessarily need to know which objects it will receive.
<i>import_to_part</i>	An array of size <i>num_import</i> listing the partitions to which objects will be imported. This array may be NULL, as the processor does not necessarily need to know from which objects it will receive.
<i>num_export</i>	The number of objects that will be sent from this processor to other processors.
<i>export_global_ids</i>	An array of <i>num_export</i> global IDs of objects to be sent from this processor.
<i>export_local_ids</i>	An array of <i>num_export</i> local IDs of objects to be sent from this processor.
<i>export_procs</i>	An array of size <i>num_export</i> listing the processor IDs of the destination processors.
<i>export_to_part</i>	An array of size <i>num_export</i> listing the partitions to which objects will be sent.
<i>ierr</i>	Error code to be set by function.

**Default:**

No processing is done if a **ZOLTAN\_MID\_MIGRATE\_PP\_FN** is not registered.

---

C and C++:     typedef void **ZOLTAN\_POST\_MIGRATE\_PP\_FN** (  
                   void \**data*,  
                   int *num\_gid\_entries*,  
                   int *num\_lid\_entries*,  
                   int *num\_import*,  
                   [ZOLTAN\\_ID\\_PTR](#) *import\_global\_ids*,  
                   [ZOLTAN\\_ID\\_PTR](#) *import\_local\_ids*,  
                   int \**import\_procs*,  
                   int \**import\_to\_part*,  
                   int *num\_export*,  
                   [ZOLTAN\\_ID\\_PTR](#) *export\_global\_ids*,  
                   [ZOLTAN\\_ID\\_PTR](#) *export\_local\_ids*,  
                   int \**export\_procs*,  
                   int \**export\_to\_part*,  
                   int \**ierr*);

FORTRAN:     SUBROUTINE **Post\_Migrate\_PP**(*data*, *num\_gid\_entries*, *num\_lid\_entries*,  
                   *num\_import*, *import\_global\_ids*, *import\_local\_ids*, *import\_procs*, *import\_to\_part*,  
                   *num\_export*, *export\_global\_ids*, *export\_local\_ids*, *export\_procs*, *export\_to\_part*, *ierr*)  
                   <*type-data*>, INTENT(INOUT) :: *data*  
                   INTEGER(Zoltan\_INT), INTENT(IN) :: *num\_gid\_entries*, *num\_lid\_entries*  
                   INTEGER(Zoltan\_INT), INTENT(IN) :: *num\_import*, *num\_export*  
                   INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *import\_global\_ids*,  
                   *export\_global\_ids*  
                   INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *import\_local\_ids*,  
                   *export\_local\_ids*  
                   INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *import\_procs*, *export\_procs*  
                   INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *import\_to\_part*,  
                   *export\_to\_part*  
                   INTEGER(Zoltan\_INT), INTENT(OUT) :: *ierr*  
  
                   <*type-data*> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or  
                   REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE),  
                   DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section  
                   on [Fortran query functions](#) for an explanation.

---

A **ZOLTAN\_POST\_MIGRATE\_PP\_FN** query function performs any post-processing desired by the application. If it is registered, it is called at the end of the [Zoltan\\_Migrate](#) routine. The arguments passed to [Zoltan\\_Migrate](#) are made available for use in the post-processing routine.

**Function Type:**     **ZOLTAN\_POST\_MIGRATE\_PP\_FN\_TYPE**

**Arguments:**

<i>data</i>	Pointer to user-defined data.
<i>num_gid_entries</i>	The number of array entries used to describe a single global ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_GID_ENTRIES</a> .



<i>num_lid_entries</i>	The number of array entries used to describe a single local ID. This value is the maximum value over all processors of the parameter <a href="#">NUM_LID_ENTRIES</a> .
<i>num_import</i>	The number of objects that will be received by this processor. An array of <i>num_import</i> global IDs of objects to be received by this processor.
<i>import_global_ids</i>	This array may be NULL, as the processor does not necessarily need to know which objects it will receive.
<i>import_local_ids</i>	An array of <i>num_import</i> local IDs of objects to be received by this processor. This array may be NULL, as the processor does not necessarily need to know which objects it will receive.
<i>import_procs</i>	An array of size <i>num_import</i> listing the processor IDs of the source processors. This array may be NULL, as the processor does not necessarily need to know which objects it will receive.
<i>import_to_part</i>	An array of size <i>num_import</i> listing the partitions to which objects will be imported. This array may be NULL, as the processor does not necessarily need to know from which objects it will receive.
<i>num_export</i>	The number of objects that will be sent from this processor to other processors.
<i>export_global_ids</i>	An array of <i>num_export</i> global IDs of objects to be sent from this processor.
<i>export_local_ids</i>	An array of <i>num_export</i> local IDs of objects to be sent from this processor.
<i>export_procs</i>	An array of size <i>num_export</i> listing the processor IDs of the destination processors.
<i>export_to_part</i>	An array of size <i>num_export</i> listing the partitions to which objects will be sent.
<i>ierr</i>	Error code to be set by function.

**Default:**

No post-processing is done if a **ZOLTAN\_POST\_MIGRATE\_PP\_FN** is not registered.

C:           typedef void **ZOLTAN\_PRE\_MIGRATE\_FN** (  
                   void \*data,  
                   int num\_gid\_entries,  
                   int num\_lid\_entries,  
                   int num\_import,  
                   [ZOLTAN\\_ID\\_PTR](#) import\_global\_ids,  
                   [ZOLTAN\\_ID\\_PTR](#) import\_local\_ids,  
                   int \*import\_procs,  
                   int num\_export,  
                   [ZOLTAN\\_ID\\_PTR](#) export\_global\_ids,  
                   [ZOLTAN\\_ID\\_PTR](#) export\_local\_ids,  
                   int \*export\_procs,  
                   int \*ierr);



**FORTTRAN:**     SUBROUTINE ***Pre\_Migrate***(*data*, *num\_gid\_entries*, *num\_lid\_entries*, *num\_import*,  
*import\_global\_ids*, *import\_local\_ids*, *import\_procs*, *num\_export*, *export\_global\_ids*,  
*export\_local\_ids*, *export\_procs*, *ierr*)  
<type-data>, INTENT(INOUT) :: *data*  
INTEGER(Zoltan\_INT), INTENT(IN) :: *num\_gid\_entries*, *num\_lid\_entries*  
INTEGER(Zoltan\_INT), INTENT(IN) :: *num\_import*, *num\_export*  
INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *import\_global\_ids*,  
*export\_global\_ids*  
INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *import\_local\_ids*,  
*export\_local\_ids*  
INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *import\_procs*, *export\_procs*  
INTEGER(Zoltan\_INT), INTENT(OUT) :: *ierr*

<type-data> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or  
REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE),  
DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section  
on [Fortran query functions](#) for an explanation.

---

A **ZOLTAN\_PRE\_MIGRATE\_FN** query function performs any pre-processing desired by applications  
using [Zoltan\\_Help\\_Migrate](#). Its function is analogous to [ZOLTAN\\_PRE\\_MIGRATE\\_PP\\_FN](#), but it  
cannot be used with [Zoltan\\_Migrate](#).

**Function Type:**     **ZOLTAN\_PRE\_MIGRATE\_FN\_TYPE**

**Arguments:**

All arguments are analogous to those in [ZOLTAN\\_PRE\\_MIGRATE\\_PP\\_FN](#).  
Partition-assignment arguments *import\_to\_part* and *export\_to\_part* are not  
included, as processor and partitions numbers are considered to be the same in  
[Zoltan\\_Help\\_Migrate](#).

**Default:**

No pre-processing is done if a **ZOLTAN\_PRE\_MIGRATE\_FN** is not registered.

---

**C:**             typedef void **ZOLTAN\_MID\_MIGRATE\_FN** (  
                  void \**data*,  
                  int *num\_gid\_entries*,  
                  int *num\_lid\_entries*,  
                  int *num\_import*,  
                  [ZOLTAN\\_ID\\_PTR](#) *import\_global\_ids*,  
                  [ZOLTAN\\_ID\\_PTR](#) *import\_local\_ids*,  
                  int \**import\_procs*,  
                  int *num\_export*,  
                  [ZOLTAN\\_ID\\_PTR](#) *export\_global\_ids*,  
                  [ZOLTAN\\_ID\\_PTR](#) *export\_local\_ids*,  
                  int \**export\_procs*,

```
int *ierr);
```

**FORTTRAN:** SUBROUTINE *Mid\_Migrate*(data, num\_gid\_entries, num\_lid\_entries, num\_import, import\_global\_ids, import\_local\_ids, import\_procs, num\_export, export\_global\_ids, export\_local\_ids, export\_procs, ierr)  
 <type-data>, INTENT(INOUT) :: data  
 INTEGER(Zoltan\_INT), INTENT(IN) :: num\_gid\_entries, num\_lid\_entries  
 INTEGER(Zoltan\_INT), INTENT(IN) :: num\_import, num\_export  
 INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: import\_global\_ids, export\_global\_ids  
 INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: import\_local\_ids, export\_local\_ids  
 INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: import\_procs, export\_procs  
 INTEGER(Zoltan\_INT), INTENT(OUT) :: ierr

<type-data> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE), DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where x is 1, 2, 3 or 4. See the section on [Fortran query functions](#) for an explanation.

---

A **ZOLTAN\_MID\_MIGRATE\_FN** query function performs any mid-migration processing desired by applications using [Zoltan\\_Help\\_Migrate](#). Its function is analogous to [ZOLTAN\\_MID\\_MIGRATE\\_PP\\_FN](#), but it cannot be used with [Zoltan\\_Migrate](#).

**Function Type:** ZOLTAN\_MID\_MIGRATE\_FN\_TYPE

**Arguments:**

All arguments are analogous to those in [ZOLTAN\\_MID\\_MIGRATE\\_PP\\_FN](#). Partition-assignment arguments *import\_to\_part* and *export\_to\_part* are not included, as processor and partitions numbers are considered to be the same in [Zoltan\\_Help\\_Migrate](#).

**Default:**

No processing is done if a **ZOLTAN\_MID\_MIGRATE\_FN** is not registered.

---

C:           typedef void **ZOLTAN\_POST\_MIGRATE\_FN** (  
               void *\*data*,  
               int *num\_gid\_entries*,  
               int *num\_lid\_entries*,  
               int *num\_import*,  
               [ZOLTAN\\_ID\\_PTR](#) *import\_global\_ids*,  
               [ZOLTAN\\_ID\\_PTR](#) *import\_local\_ids*,  
               int *\*import\_procs*,  
               int *num\_export*,  
               [ZOLTAN\\_ID\\_PTR](#) *export\_global\_ids*,  
               [ZOLTAN\\_ID\\_PTR](#) *export\_local\_ids*,  
               int *\*export\_procs*,  
               int *\*ierr*);

FORTTRAN: SUBROUTINE **Post\_Migrate**(*data*, *num\_gid\_entries*, *num\_lid\_entries*, *num\_import*,  
*import\_global\_ids*, *import\_local\_ids*, *import\_procs*, *num\_export*, *export\_global\_ids*,  
*export\_local\_ids*, *export\_procs*, *ierr*)  
<type-data>, INTENT(INOUT) :: *data*  
INTEGER(Zoltan\_INT), INTENT(IN) :: *num\_gid\_entries*, *num\_lid\_entries*  
INTEGER(Zoltan\_INT), INTENT(IN) :: *num\_import*, *num\_export*  
INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *import\_global\_ids*,  
*export\_global\_ids*  
INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *import\_local\_ids*,  
*export\_local\_ids*  
INTEGER(Zoltan\_INT), INTENT(IN), DIMENSION(\*) :: *import\_procs*, *export\_procs*  
INTEGER(Zoltan\_INT), INTENT(OUT) :: *ierr*

<type-data> can be any of INTEGER(Zoltan\_INT), DIMENSION(\*) or  
REAL(Zoltan\_FLOAT), DIMENSION(\*) or REAL(Zoltan\_DOUBLE),  
DIMENSION(\*) or TYPE(Zoltan\_User\_Data\_x) where *x* is 1, 2, 3 or 4. See the section  
on [Fortran query functions](#) for an explanation.

---

A **ZOLTAN\_POST\_MIGRATE\_FN** query function performs any post-processing desired by  
applications using [Zoltan\\_Help\\_Migrate](#). Its function is analogous to  
[ZOLTAN\\_POST\\_MIGRATE\\_PP\\_FN](#), but it cannot be used with [Zoltan\\_Migrate](#).

**Function Type:**    **ZOLTAN\_POST\_MIGRATE\_FN\_TYPE**

**Arguments:**

All arguments are analogous to those in [ZOLTAN\\_POST\\_MIGRATE\\_PP\\_FN](#).  
Partition-assignment arguments *import\_to\_part* and *export\_to\_part* are not  
included, as processor and partitions numbers are considered to be the same in  
[Zoltan\\_Help\\_Migrate](#).

**Default:**

No post-processing is done if a **ZOLTAN\_POST\_MIGRATE\_FN** is not  
registered.

---

[\[Table of Contents](#) | [Next: Zoltan Parameters and Output Levels](#) | [Previous: Load-Balancing Query Functions\]](#)

# Zoltan Parameters and Output Levels

The behavior of Zoltan is controlled by several [parameters](#) and [debugging-output levels](#). These parameters can be set by calls to [Zoltan\\_Set\\_Param](#). Reasonable [default values](#) for all parameters are specified by Zoltan. Many of the parameters are specific to individual algorithms, and are listed in the descriptions of those algorithms. However, the parameters below have meaning across the entire library.

## General Parameters

The following parameters apply to the entire Zoltan library. While reasonable [default values](#) for all parameters are specified by Zoltan, applications can change these values through calls to [Zoltan\\_Set\\_Param](#).

### Parameters:

<a href="#"><i>NUM_GID_ENTRIES</i></a>	The number of unsigned integers that should be used to represent a global identifier (ID). Values greater than zero are accepted.
<a href="#"><i>NUM_LID_ENTRIES</i></a>	The number of unsigned integers that should be used to represent a local identifier (ID). Values greater than or equal to zero are accepted.
<a href="#"><i>DEBUG_LEVEL</i></a>	An integer indicating how much debugging information is printed by Zoltan. Higher values of <a href="#"><i>DEBUG_LEVEL</i></a> produce more output and potentially slow down Zoltan's computations. The least output is produced when <a href="#"><i>DEBUG_LEVEL</i></a> = 0. <a href="#"><i>DEBUG_LEVEL</i></a> primarily controls Zoltan's behavior; most algorithms have their own parameters to control their output level. Values used within Zoltan are listed <a href="#">below</a> . Note: Because some debugging levels use processor synchronization, all processors should use the same value of <a href="#"><i>DEBUG_LEVEL</i></a> .
<i>DEBUG_PROCESSOR</i>	Processor number from which trace output should be printed when <a href="#"><i>DEBUG_LEVEL</i></a> is 5.
<i>DEBUG_MEMORY</i>	Integer indicating the amount of low-level debugging information about memory-allocation should be kept by Zoltan's <a href="#">Memory Management utilities</a> . Valid values are 0, 1, 2, and 3.
<i>OBJ_WEIGHT_DIM</i>	The number of weights associated with an object. If this parameter is zero, all objects have equal weight. Some algorithms may not support multiple (multidimensional) weights.
<i>EDGE_WEIGHT_DIM</i>	The number of weights associated with an edge. If this parameter is zero, all edges have equal weight. Many algorithms do not support multiple (multidimensional) weights.

<i>TIMER</i>	The timer with which you wish to measure time. Valid choices are <i>wall</i> (based on <i>MPI_Wtime</i> ), <i>cpu</i> (based on the ANSI C library function <i>clock</i> ), and <i>user</i> . The resolution may be poor, as low as 1/60th of a second, depending upon your platform.
<i>USE_MACHINE_DESC</i>	Currently unused; will be used when heterogeneous computers are supported.
<i>MACHINE_DESC_FILE</i>	Currently unused; will be used when heterogeneous computers are supported.

**Default Values:**

*NUM\_GID\_ENTRIES* = 1  
*NUM\_LID\_ENTRIES* = 1  
*DEBUG\_LEVEL* = 1  
*DEBUG\_PROCESSOR* = 0  
*DEBUG\_MEMORY* = 1  
*OBJ\_WEIGHT\_DIM* = 0  
*EDGE\_WEIGHT\_DIM* = 0  
*TIMER* = wall  
*USE\_MACHINE\_DESC* = 0  
*MACHINE\_DESC\_FILE* = /etc/local/Zoltan\_Machine\_Desc

---

## Debugging Levels in Zoltan

The *DEBUG\_LEVEL* parameter determines how much debugging information is printed to *stdout* by Zoltan. It is set by a call to [Zoltan\\_Set\\_Param](#). Higher values of *DEBUG\_LEVEL* produce more output and can slow down Zoltan's computations, especially when the output is printed by one processor at a time. The least output is produced when *DEBUG\_LEVEL* = 0.

Descriptions of the output produced by Zoltan for each value of *DEBUG\_LEVEL* are included below. For a given *DEBUG\_LEVEL* value *n*, all output for values less than or equal to *n* is produced.

Some high debugging levels use processor synchronization to force processors to write one-at-a-time. For example, when *DEBUG\_LEVEL* is greater than or equal to eight, each processor writes its list in turn so that the lists from all processors are not jumbled together in the output. This synchronization requires all processors to use the same value of *DEBUG\_LEVEL*.

### *DEBUG\_LEVEL* Output Produced

0	Quiet mode; no output unless an error or warning is produced.
1	Values of all parameters set by <a href="#">Zoltan_Set_Param</a> and used by Zoltan.
2	Timing information for Zoltan's main routines.
3	Timing information within Zoltan's algorithms (support by algorithms is optional).
4	
5	Trace information (enter/exit) for major Zoltan interface routines (printed by the processor specified by the <a href="#">DEBUG_PROCESSOR</a> parameter).

- 6 Trace information (enter/exit) for major Zoltan interface routines (printed by all processors).
- 7 More detailed trace information in major Zoltan interface routines.
- 8 List of objects to be imported to and exported from each processor. <sup>1</sup>
- 9
- 10 Maximum debug output; may include algorithm-specific output. <sup>1</sup>

<sup>1</sup> *Output may be serialized; that is, one processor may have to complete its output before the next processor is allowed to begin its output. This serialization is not scalable and can significantly increase execution time on large number of processors.*

---

[\[Table of Contents\]](#) | [\[Next: Load-Balancing Algorithms\]](#) | [\[Previous: Migration Query Functions\]](#)



# Load-Balancing Algorithms

The following dynamic load-balancing algorithms are currently included in the Zoltan library:

[Recursive Coordinate Bisection \(RCB\)](#)

[Recursive Inertial Bisection \(RIB\)](#)

[Hilbert Space-Filling Curve \(HSFC\)](#)

[Refinement Tree Based Partitioning \(REFTREE\)](#)

[Graph partitioning: ParMETIS \(GRAPH or PARMETIS\)](#)

[Graph partitioning: Jostle \(JOSTLE\)](#)

[Hypergraph partitioning: PHG \(HYPERGRAPH or PHG\)](#)

[Octree Partitioning \(OCTPART\)](#)

The parenthetical string is the parameter value for [LB\\_METHOD](#) parameter; the parameter is set through a call to [Zoltan\\_Set\\_Param](#).

For further analysis and discussion of the algorithms, see [[Hendrickson and Devine](#)].

---

## Load-Balancing Parameters

While the overall behavior of Zoltan is controlled by [general Zoltan parameters](#), the behavior of each load-balancing method is controlled by parameters specific to partitioning which are also set by calls to [Zoltan\\_Set\\_Param](#). Many of these parameters are specific to individual partitioning algorithms, and are listed in the descriptions of the individual algorithms. However, several have meaning across multiple partitioning algorithms. These load-balancing parameters are described below. Unless indicated otherwise, these parameters apply to both [Zoltan\\_LB\\_Partition](#) and [Zoltan\\_LB\\_Balance](#).

If any processor sets this parameter, NUM\_LOCAL\_PARTITIONS is assumed to be zero on processors not setting this parameter.

### Parameters:

*LB\_METHOD*

The load-balancing algorithm used by Zoltan is specified by this parameter. Valid values are

"RCB" (for [recursive coordinate bisection](#)),  
 "RIB" (for [recursive inertial bisection](#)),  
 "HSFC" (for [Hilbert space-filling curve partitioning](#)),  
 "GRAPH" or "PARMETIS" (for any of the methods in the [ParMETIS](#) library),  
 "JOSTLE" (for any of the methods in the [Jostle](#) library),  
 "HYPERGRAPH" or "PHG" (for [hypergraph partitioning](#)),  
 "OCTPART" (for [octree partitioning](#)),  
 "REFTREE" (for [refinement tree based partitioning](#)), and  
 "NONE" (for no load-balancing).

*NUM\_GLOBAL\_PARTITIONS*

The total number of partitions to be generated by a call to [Zoltan\\_LB\\_Partition](#). Integer values greater than zero are accepted. Not valid for [Zoltan\\_LB\\_Balance](#).

*NUM\_LOCAL\_PARTITIONS*

The number of partitions to be generated on this processor by a call to [Zoltan\\_LB\\_Partition](#). Integer values greater than or equal to zero are accepted. Not valid for [Zoltan\\_LB\\_Balance](#).

*RETURN\_LISTS*

The lists returned by calls to [Zoltan\\_LB\\_Partition](#) or [Zoltan\\_LB\\_Balance](#). Valid values are

"IMPORT", to return only information about objects to be imported to a processor  
 "EXPORT", to return only information about objects to be exported from a processor  
 "ALL", or "IMPORT AND EXPORT" (or any string with both "IMPORT" and "EXPORT" in it) to return both import and export information  
 "PARTITION ASSIGNMENTS" (or any string with "PARTITION" in it) to return the new process and partition assignment of every local object, including those not being exported.  
 "NONE", to return neither import nor export information

*REMAP*

Within [Zoltan\\_LB\\_Partition](#) or [Zoltan\\_LB\\_Balance](#), renumber partitions to maximize overlap between the old decomposition and the new decomposition (to reduce data movement from old to new decompositions). Valid values are "0" (no remapping) or "1" (remapping). Requests for remapping are ignored when, in the new decomposition, a partition is spread across multiple processors or partition sizes are specified using [Zoltan\\_LB\\_Set\\_Part\\_Sizes](#).

*IMBALANCE\_TOL*

The amount of load imbalance the partitioning algorithm should deem acceptable. The load on each processor is computed as the sum of the weights of objects it is assigned. The imbalance is then computed as the maximum load divided by the average load. An value for *IMBALANCE\_TOL* of 1.2 indicates that 20% imbalance is OK; that is, the maximum over the average shouldn't exceed 1.2.

*MIGRATE\_ONLY\_PROC\_CHANGES* If this value is set to TRUE (non-zero), Zoltan's migration functions will migrate only objects moving to new processors. They will not migrate objects for which only the partition number has changed; the objects' processor numbers must change as well. If this value is set to FALSE (zero), Zoltan's migration functions will migrate all objects with new partition or processor assignments.

*AUTO\_MIGRATE*

If this value is set to TRUE (non-zero), Zoltan will automatically perform the data migration during calls to [Zoltan\\_LB\\_Partition](#) or [Zoltan\\_LB\\_Balance](#). A full discussion of automatic migration can be found in the description of the [migration interface functions](#).

**Default Values:**

*LB\_METHOD* = RCB

*NUM\_GLOBAL\_PARTITIONS* = Number of processors specified in [Zoltan\\_Create](#).

*NUM\_LOCAL\_PARTITIONS* = 1

*RETURN\_LISTS* = ALL

*REMAP* = 1

*IMBALANCE\_TOL* = 1.1

*MIGRATE\_ONLY\_PROC\_CHANGES* = 1

*AUTO\_MIGRATE* = FALSE

---

[\[Table of Contents\]](#) | [Next: Recursive Coordinate Bisection \(RCB\)](#) | [Previous: Zoltan Parameters and Output Levels](#)

# Recursive Coordinate Bisection (RCB)

An implementation of Recursive Coordinate Bisection (RCB) due to Steve Plimpton of Sandia National Laboratories is included in Zoltan. RCB was first proposed as a static load-balancing algorithm by [Berger and Bokhari](#), but is attractive as a dynamic load-balancing algorithm because it implicitly produces incremental partitions. In RCB, the computational domain is first divided into two regions by a cutting plane orthogonal to one of the coordinate axes so that half the work load is in each of the sub-regions. The splitting direction is determined by computing in which coordinate direction the set of objects is most elongated, based upon the geometric locations of the objects. The sub-regions are then further divided by recursive application of the same splitting algorithm until the number of sub-regions equals the number of processors. Although this algorithm was first devised to cut into a number of sets which is a power of two, the set sizes in a particular cut needn't be equal. By adjusting the partition sizes appropriately, any number of equally-sized sets can be created. If the parallel machine has processors with different speeds, sets with nonuniform sizes can also be easily generated. The Zoltan implementation of RCB has several parameters which can be modified by the [Zoltan\\_Set\\_Param](#) function. A recent feature is that RCB allows multiple weights; that is, one can balance with respect to several load criteria simultaneously. Note that there is no guarantee that a desired load balance tolerance can be achieved using RCB, especially in the multiconstraint case.

Information about the sub-regions generated by RCB can be obtained by an application through calls to [Zoltan\\_RCB\\_Box](#). This function is not required to perform load balancing; it only provides auxiliary information to an application.

## Method String:

## RCB

### Parameters:

*RCB\_OVERALLOC*

The amount by which to over-allocate temporary storage arrays for objects within the RCB algorithm when additional storage is due to changes in processor assignments.

1.0 = no extra storage allocated; 1.5 = 50% extra storage; etc.

*RCB\_REUSE*

Flag to indicate whether to use previous cuts as initial guesses for the current RCB invocation.

0 = don't use previous cuts; 1 = use previous cuts.

*RCB\_OUTPUT\_LEVEL*

Flag controlling the amount of timing and diagnostic output the routine produces.

0 = no output; 1 = print summary; 2 = print data for each processor.

*CHECK\_GEOM*

Flag controlling the invocation of input and output error checking.

0 = don't do checking; 1 = do checking.

*KEEP\_CUTS*

Should information about the cuts determining the RCB decomposition be retained? It costs a bit of time to do so, but this information is necessary if application wants to add more objects to the decomposition via calls to [Zoltan\\_LB\\_Point\\_PP\\_Assign](#) or to [Zoltan\\_LB\\_Box\\_PP\\_Assign](#).

0 = don't keep cuts; 1 = keep cuts.

*AVERAGE\_CUTS*

When set to one, coordinates of RCB cutting planes are computed to be the average of the coordinates of the closest object on each side of the cut. Otherwise, coordinates of cutting planes may equal those of one of the closest objects.

0 = don't average cuts; 1 = average cuts.

*RCB\_LOCK\_DIRECTIONS*

Flag that determines whether the order of the directions of the cuts is kept constant after they are determined the first time RCB is called.

0 = don't lock directions; 1 = lock directions.

*RCB\_SET\_DIRECTIONS*

If this flag is set, the order of cuts is changed so that all of the cuts in any direction are done as a group. The number of cuts in each direction is determined and then the value of the parameter is used to determine the order that those cuts are made in. When 1D and 2D problems are partitioned, the directions corresponding to unused dimensions are ignored.

0 = don't order cuts; 1 = xyz; 2 = xzy; 3 = yzx; 4 = yxz; 5 = zxy; 6 = zyx;

*RCB\_RECTILINEAR\_BLOCKS*

Flag controlling the shape of the resulting regions. If this option is specified, then when a cut is made, all of the dots located on the cut are moved to the same side of the cut. The resulting regions are then rectilinear. When these dots are treated as a group, then the resulting load balance may not be as good as when the group of dots is split by the cut.

0 = move dots individually; 1 = move dots in groups.

*REDUCE\_DIMENSIONS*

When a 3 dimensional geometry is almost flat, it may make more sense to treat it as a 2 dimensional geometry when applying the RCB algorithm. In this case, a 2 dimensional RCB calculation is applied to a plane that corresponds with the geometry. (This results in cuts that, while still orthogonal, may no longer be axis aligned.) If this parameter is set to **1**, a 3 dimensional geometry will be treated as 2 dimensional if it is very flat, or 1 dimensional if it is very thin. A 2 dimensional geometry will be treated as 1 dimensional if it is very thin.

*DEGENERATE\_RATIO*

If the **REDUCE\_DIMENSIONS** parameter is set, then this parameter determines when a geometry is considered to be degenerate. A bounding box which is oriented to the geometry is constructed, and the lengths of its sides are tested against a ratio of 1 : **DEGENERATE\_RATIO**.

***RCB\_RECOMPUTE\_BOX***

Flag indicating whether the bounding box of set of partitions is recomputed at each level of recursion. By default, the longest direction of the bounding box is cut during bisection. Recomputing the bounding box at each level of recursion can produce more effective cut directions for unusually shaped geometries; the computation does, however, take additional time and communication, and may cause cut directions to vary from one invocation of RCB to the next.

0 = don't recompute the bounding box; 1 = recompute the box.

***OBJ\_WEIGHTS\_COMPARABLE***

In the multiconstraint case, are the object weights comparable? Do they have the same units and is the scaling meaningful? For example, if the *j*th weight corresponds to the expected time in phase *j* (measured in seconds), set this parameter to 1. (0 = incomparable, 1 = comparable)

***RCB\_MULTICRITERIA\_NORM***

Norm used in multicriteria algorithm; this determines how to balance the different weight constraints. Valid values are 1,2, and 3. Roughly, if the weights correspond to different phases, then the value 1 (1-norm) tries to minimize the total time (sum over all phases) while the value 3 (max-norm) attempts to minimize the worst imbalance in any phase. The 2-norm does something in between. Try a different value if you're not happy with the balance.

***RCB\_MAX\_ASPECT\_RATIO***

Maximum allowed ratio between the largest and smallest side of a subdomain. Must be > 1.

**Default:**

*RCB\_OVERALLOC* = 1.0

*RCB\_REUSE* = 0

*RCB\_OUTPUT\_LEVEL* = 0

*CHECK\_GEOM* = 1

*KEEP\_CUTS* = 0

*AVERAGE\_CUTS* = 0

*RCB\_LOCK\_DIRECTIONS* = 0

*REDUCE\_DIMENSIONS* = 0

*DEGENERATE\_RATIO* = 10

*RCB\_SET\_DIRECTIONS* = 0

*RCB\_RECTILINEAR\_BLOCKS* = 0

*RCB\_RECOMPUTE\_BOX* = 0

*OBJ\_WEIGHTS\_COMPARABLE* = 0

*RCB\_MULTICRITERIA\_NORM* = 1

*RCB\_MAX\_ASPECT\_RATIO* = 10

**Required Query Functions:**

[ZOLTAN\\_NUM\\_OBJ\\_FN](#)

[ZOLTAN\\_OBJ\\_LIST\\_FN](#) or

[ZOLTAN\\_FIRST\\_OBJ\\_FN](#)/[ZOLTAN\\_NEXT\\_OBJ\\_FN](#) pair



[ZOLTAN\\_NUM\\_GEOM\\_FN](#)[ZOLTAN\\_GEOM\\_MULTI\\_FN](#) or [ZOLTAN\\_GEOM\\_FN](#)


---

```

C:      int Zoltan_RCB_Box (
        struct Zoltan_Struct * zz,
        int part,
        int *ndim,
        double *xmin,
        double *ymin,
        double *zmin,
        double *xmax,
        double *ymax,
        double *zmax);

FORTRAN:  FUNCTION Zoltan_RCB_Box(zz, part,ndim, xmin, ymin, zmin, xmax, ymax, zmax)
          INTEGER(Zoltan_INT) :: Zoltan_RCB_Box
          TYPE(Zoltan_Struct), INTENT(IN) :: zz
          INTEGER(Zoltan_INT), INTENT(IN) :: part
          INTEGER(Zoltan_INT), INTENT(OUT) :: ndim
          REAL(Zoltan_DOUBLE), INTENT(OUT) :: xmin, ymin, zmin, xmax, ymax, zmax

```

---

In many settings, it is useful to know a partition's bounding box generated by RCB. This bounding box describes the region of space assigned to a given partition. Given an RCB decomposition of space and a partition number, **Zoltan\_RCB\_Box** returns the lower and upper corners of the region of space assigned to the partition. To use this routine, the parameter **KEEP\_CUTS** must be set to TRUE when the decomposition is generated. This parameter will cause the sequence of geometric cuts to be saved, which is necessary for **Zoltan\_RCB\_Box** to do its job.

**Arguments:**

<i>zz</i>	Pointer to the Zoltan structure created by <a href="#"><u>Zoltan_Create</u></a> .
<i>part</i>	Partition number of partition for which the bounding box should be returned.
<i>ndim</i>	Upon return, the number of dimensions in the partitioned geometry.
<i>xmin</i> , <i>ymin</i> , <i>zmin</i>	Upon return, the coordinates of the lower extent of bounding box for the partition. If the geometry is two-dimensional, <i>zmin</i> is -DBL_MAX. If the geometry is one-dimensional, <i>ymin</i> is -DBL_MAX.
<i>xmax</i> , <i>ymax</i> , <i>zmax</i>	Upon return, the coordinates of the upper extent of bounding box for the partition. If the geometry is two-dimensional, <i>zmax</i> is DBL_MAX. If the geometry is one-dimensional, <i>ymax</i> is DBL_MAX.

**Returned Value:**

int      [Error code](#).

---

[[Table of Contents](#) | [Next: Recursive Inertial Bisection \(RIB\)](#) | [Previous: Load-Balancing Algorithms](#)]



# Recursive Inertial Bisection (RIB)

An implementation of Recursive Inertial Bisection (RIB) is included in Zoltan. RIB was proposed as a load-balancing algorithm by [Williams](#) and later studied by [Taylor and Nour-Omid](#), but its origin is unclear. RIB is similar to RCB in that it divides the domain based on the location of the objects being partitioned by use of cutting planes. In RIB, the computational domain is first divided into two regions by a cutting plane orthogonal to the longest direction of the domain so that half the work load is in each of the sub-regions. The sub-regions are then further divided by recursive application of the same splitting algorithm until the number of sub-regions equals the number of processors. Although this algorithm was first devised to cut into a number of sets which is a power of two, the set sizes in a particular cut needn't be equal. By adjusting the partition sizes appropriately, any number of equally-sized sets can be created. If the parallel machine has processors with different speeds, sets with nonuniform sizes can also be easily generated. The Zoltan implementation of RIB has several parameters which can be modified by the [Zoltan\\_Set\\_Param](#) function.

## Method String:

**RIB**

## Parameters:

<i>RIB_OVERALLOC</i>	The amount by which to over-allocate temporary storage arrays for objects within the RIB algorithm when additional storage is due to changes in processor assignments. 1.0 = no extra storage allocated; 1.5 = 50% extra storage; etc.
<i>RIB_OUTPUT_LEVEL</i>	Flag controlling the amount of timing and diagnostic output the routine produces. 0 = no output; 1 = print summary; 2 = print data for each processor.
<i>CHECK_GEOM</i>	Flag controlling the invocation of input and output error checking. 0 = don't do checking; 1 = do checking.
<i>KEEP_CUTS</i>	Should information about the cuts determining the RIB decomposition be retained? It costs a bit of time to do so, but this information is necessary if application wants to add more objects to the decomposition via calls to <a href="#">Zoltan_LB_Point_PP_Assign</a> or to <a href="#">Zoltan_LB_Box_PP_Assign</a> . 0 = don't keep cuts; 1 = keep cuts.
<i>AVERAGE_CUTS</i>	When set to one, coordinates of RIB cutting planes are computed to be the average of the coordinates of the closest object on each side of the cut. Otherwise, coordinates of cutting planes may equal those of one of the closest objects. 0 = don't average cuts; 1 = average cuts.

When a 3 dimensional geometry is almost flat, it may make more sense to *REDUCE\_DIMENSIONS* treat it as a 2 dimensional geometry when applying the RIB algorithm. (Coordinate values in the omitted direction are ignored for the purposes of partitioning.) If this parameter is set to **1**, a 3 dimensional geometry will be treated as 2 dimensional if it is very flat, or 1 dimensional if it is very thin. A 2 dimensional geometry will be treated as 1 dimensional if it is very thin. If the **REDUCE\_DIMENSIONS** parameter is set, then this parameter *DEGENERATE\_RATIO* determines when a geometry is considered to be degenerate. A bounding box which is oriented to the geometry is constructed, and the lengths of its sides are tested against a ratio of 1 : **DEGENERATE\_RATIO**.

**Default:**

*RIB\_OVERALLOC* = 1.0  
*RIB\_OUTPUT\_LEVEL* = 0  
*CHECK\_GEOM* = 1  
*KEEP\_CUTS* = 0  
*AVERAGE\_CUTS* = 0  
*REDUCE\_DIMENSIONS* = 0  
*DEGENERATE\_RATIO* = 10

**Required Query Functions:**

[ZOLTAN\\_NUM\\_OBJ\\_FN](#)  
[ZOLTAN\\_OBJ\\_LIST\\_FN](#) or  
[ZOLTAN\\_FIRST\\_OBJ\\_FN](#)/[ZOLTAN\\_NEXT\\_OBJ\\_FN](#) pair  
[ZOLTAN\\_NUM\\_GEOM\\_FN](#)  
[ZOLTAN\\_GEOM\\_MULTI\\_FN](#) or [ZOLTAN\\_GEOM\\_FN](#)

---

[\[Table of Contents](#) | [Next: Hilbert Space-Filling Curve Partitioning](#) | [Previous: Recursive Coordinate Bisection \(RCB\)\]](#)

# Hilbert Space Filling Curve (HSFC)

The Inverse Hilbert Space-Filling Curve functions map a point in one, two or three dimensions into the interval  $[0,1]$ . The Hilbert functions that map  $[0, 1]$  to normal spatial coordinates are also provided. (The one-dimensional inverse Hilbert curve is defined here as the identity function,  $f(x)=x$  for all  $x$ .)

The HSFC partitioning algorithm seeks to divide  $[0,1]$  into  $P$  intervals each containing the same weight of objects associated to these intervals by their inverse Hilbert coordinates.  $N$  bins are created (where  $N > P$ ) to partition  $[0,1]$ . The weights in each bin are summed across all processors. A greedy algorithm sums the bins (from left to right) placing a cut when the desired weight for current partition interval is achieved. This process is repeated as needed to improve partitioning tolerance by a technique that maintains the same total number of bins but refines the bins previously containing a cut.

HSFC returns an warning if the final imbalance exceeds the user specified tolerance.

This code implements both the point assign and box assign functionality. The point assign determines an appropriate partition (associated with a specific group of processors) for a new point. The box assign determines the list of processors whose associated subdomains intersect the given box. In order to use either of these routines, the user parameter `KEEP_CUTS` must be turned on. Both point assign and box assign now work for points or boxes anywhere in space even if they are exterior to the original bounding box. If a partition is empty (due to the partition being assigned zero work), it is not included in the list of partitions returned by box assign. Note: the original box assign algorithm was not rigorous and may have missed partitions. This version is both rigorous and fast.

The Zoltan implementation of HSFC has one parameter that can be modified by the [Zoltan\\_Set\\_Param](#) function.

This partitioning algorithm is loosely based on the 2D & 3D Hilbert tables used in the Octree partitioner and on the BSFC partitioning implementation by Andrew C. Bauer, Department of Engineering, State University of New York at Buffalo, as his summer project at SNL in 2001. The box assign algorithm is loosely based on the papers by Lawder referenced both in the developers guide and the code itself.

NOTE: This code can be trivially extended to any space filling curve by providing the tables implementing the curve's state transition diagram. The only dependance on the curve is through the tables and the box assign algorithm will work for all space filling curves (if we have their tables.)

Please refer to the Zoltan Developers Guide, [Appendix: Hilbert Space Filling Curve \(HSFC\)](#) for more detailed information about these algorithms.

**Method String:** HSFC  
**Parameters:**

***KEEP\_CUTS***

Information about cuts and bounding box is necessary if the application wants to add more objects to the decomposition via calls to

[Zoltan\\_LB\\_Point\\_PP\\_Assign](#) or to [Zoltan\\_LB\\_Box\\_PP\\_Assign](#).

0 = don't keep cuts; 1 = keep cuts.

***REDUCE\_DIMENSIONS***

When a 3 dimensional geometry is almost flat, it may make more sense to treat it as a 2 dimensional geometry when applying the HSFC algorithm. (Coordinate values in the omitted direction are ignored for the purposes of partitioning.) If this parameter is set to **1**, a 3 dimensional geometry will be treated as 2 dimensional if is very flat, or 1 dimensional if it very thin. And a 2 dimensional geometry will be treated as 1 dimensional if it is very thin. Turning this parameter on removes the possibility that disconnected partitions will appear on the surface of a flat 3 dimensional object.

***DEGENERATE\_RATIO***

If the **REDUCE\_DIMENSIONS** parameter is set, then this parameter determines when a geometry is considered to be flat. A bounding box which is oriented to the geometry is constructed, and the lengths of it's sides are tested against a ratio of 1 : **DEGENERATE\_RATIO**.

**Default:**

*KEEP\_CUTS* = 0

*REDUCE\_DIMENSIONS* = 0

*DEGENERATE\_RATIO* = 10

**Required Query Functions:**

[ZOLTAN\\_NUM\\_OBJ\\_FN](#)

[ZOLTAN\\_OBJ\\_LIST\\_FN](#) or

[ZOLTAN\\_FIRST\\_OBJ\\_FN](#)/[ZOLTAN\\_NEXT\\_OBJ\\_FN](#) pair

[ZOLTAN\\_NUM\\_GEOM\\_FN](#)

[ZOLTAN\\_GEOM\\_MULTI\\_FN](#) or [ZOLTAN\\_GEOM\\_FN](#)

---

[\[Table of Contents\]](#) | [\[Next: Refinement Tree Partitioning\]](#) | [\[Previous: Recursive Inertial Bisection\]](#)

# Refinement Tree Partitioning (REFTREE)

The refinement tree based partitioning method is due to William Mitchell of the National Institute of Standards and Technology [[Mitchell](#)]. It is closely related to the Octree and Space-Filling Curve methods, except it uses the tree that represents the adaptive refinement process that created the grid. This tree is constructed through the tree-based query functions.

Each node of the refinement tree corresponds to an element that occurred during the grid refinement process. The first level of the tree (the children of the root of the tree) corresponds to the initial coarse grid, one tree node per initial element. It is assumed that the initial coarse grid does not change through the execution of the program, except that the local IDs, assignment of elements to processors, and weights can change. If any other aspect of the coarse grid changes, then the Zoltan structure should be destroyed and recreated. The children of a node in the tree correspond to the elements that were created when the corresponding element was refined. The children are ordered such that a traversal of the tree creates a space-filling curve within each initial element. If the initial elements can be ordered with a contiguous path through them, then the traversal creates a space-filling curve through all the elements. Each element has a designated "in" vertex and "out" vertex, with the out vertex of one element being the same as the in vertex of the next element in the path, in other words the path goes through a vertex to move from one element to the next (and does not go out the same vertex it came in).

The user may allow Zoltan to determine the order of the coarse grid elements, or may specify the order, which might be faster or produce a better path. If Zoltan determines the order, the user can select between an order that will produce connected partitions, an order based on a Hilbert Space Filling Curve, or an order based on a Sierpinski Space Filling Curve. See the parameter REFTREE\_INITPATH below. If the user provides the order, then the in/out vertices must also be supplied. Similarly, the user may specify the order and in/out vertices of the child elements, or allow Zoltan to determine them. If the user knows how to provide a good ordering for the children, this may be significantly faster than the default general algorithm. However, accelerated forms of the ordering algorithm are provided for certain types of refinement schemes and should be used in those cases. See [ZOLTAN\\_CHILD\\_LIST\\_FN](#). If the user always specifies the order, then the vertices and in/out vertices are not used and do not have to be provided.

Weights are assigned to the nodes of the tree. These weights need not be only on the leaves (the elements of the final grid), but can also be on interior nodes (for example, to represent work on coarse grids of a multigrid algorithm). The default weights are 1.0 at the leaves and 0.0 at the interior nodes, which produces a partition based on the number of elements in each partition. An initial tree traversal is used to sum the weights, and a second traversal to cut the space-filling curve into appropriately-sized pieces and assign elements to partitions. The number of partitions is not necessarily equal to the number of processors.

The following limitations should be removed in the future.

- For multicomponent weights, only the first component is used.
- Heterogeneous architectures are not supported, in the sense that the computational load is equally divided over the processors. A vector of relative partition sizes is used to determine the weight assigned to each partition, but they are currently all equal. In the future they should be input to reflect heterogeneity.

**Method String:**        **REFTREE**

**Parameters:**

*REFTREE\_HASH\_SIZE* The size of the hash table to map from global IDs to refinement tree nodes. Larger values require more memory but may reduce search time.

**Default:**

*REFTREE\_HASH\_SIZE* = 16384

*REFTREE\_INITPATH* Determines the method for finding an order of the elements in the initial grid. "SIERPINSKI" uses a Sierpinski Space Filling Curve and is most appropriate for grids consisting of triangles. It is currently limited to 2D. "HILBERT" uses a Hilbert Space Filling Curve and is most appropriate for grids consisting of quadrilaterals or hexahedra. "CONNECTED" attempts to produce connected partitions (guaranteed for triangles and tetrahedra), however they tend to be stringy, i.e., less compact than the SFC methods. It is most appropriate when connected partitions are required. An invalid character string will invoke the default method.

**Default:**

*REFTREE\_INITPATH* = "SIERPINSKI" if the grid contains only triangles

*REFTREE\_INITPATH* = "HILBERT" otherwise

*NOTE:* In Zoltan versions 1.53 and earlier the default was "CONNECTED". To reproduce old results, use *REFTREE\_INITPATH* = "CONNECTED".

**Required Query Functions:**

[ZOLTAN\\_NUM\\_COARSE\\_OBJ\\_FN](#)

[ZOLTAN\\_COARSE\\_OBJ\\_LIST\\_FN](#) or

[ZOLTAN\\_FIRST\\_COARSE\\_OBJ\\_FN](#)/[ZOLTAN\\_NEXT\\_COARSE\\_OBJ\\_FN](#)

pair

[ZOLTAN\\_NUM\\_CHILD\\_FN](#)

[ZOLTAN\\_CHILD\\_LIST\\_FN](#)

[ZOLTAN\\_CHILD\\_WEIGHT\\_FN](#)

The following functions are needed only if the order of the initial elements will be determined by a space filling curve method:

[ZOLTAN\\_NUM\\_GEOM\\_FN](#)

[ZOLTAN\\_GEOM\\_FN](#) or [ZOLTAN\\_GEOM\\_MULTI\\_FN](#)

---

[[Table of Contents](#) | [Next: ParMETIS](#) | [Previous: Hilbert Space-Filling Curve Partitioning](#)]



Note: See also [hypergraph partitioning](#).

## Graph partitioning: ParMETIS

[ParMETIS](#) is a parallel library for graph partitioning (for static load balancing) and repartitioning (for dynamic load balancing) developed at the University of Minnesota by Karypis, Schloegel and Kumar [[Karypis and Kumar](#)]. ParMETIS is therefore strictly speaking not a method but rather a collection of methods. In the Zoltan context, ParMETIS is a method with many sub-methods. Zoltan provides an interface to all the ParMETIS (sub-)methods. The user selects which ParMETIS method to use through the parameter `PARMETIS_METHOD`. Most of the ParMETIS methods are based on either multilevel Kernighan-Lin partitioning or a diffusion algorithm. The names of the ParMETIS methods used by Zoltan are identical to those in the ParMETIS library. For further information about the various [ParMETIS](#) methods and parameters, please consult the [ParMETIS](#) User's Guide.

Graph partitioning is a useful abstraction for load balancing. The main idea is to represent the computational application as a weighted graph. The nodes or vertices in the graph correspond to objects in Zoltan. Each object may have a weight that normally represents the amount of computation. The edges or arcs in the graph usually correspond to communication costs. In graph partitioning, the problem is to find a partitioning of the graph (that is, each vertex is assigned to one out of  $k$  possible sets called partitions) that minimizes the cut size (weight) subject to the partitions having approximately equal size (weight). In repartitioning, it is assumed that a partitioning already exists. The problem is to find a good partitioning that is also "similar" in some sense to the existing partitioning. This keeps the migration cost low. All the problems described above are NP-hard so no efficient exact algorithm is known. We remark that in Zoltan 1.\*, the number of partitions is always the same as the number of MPI processes (which is normally equal to the number of processors).

We give only a brief summary of the various ParMETIS methods here; for more details see the [ParMETIS](#) documentation. The methods fall into three categories:

1. Part\* - Perform graph partitioning without consideration of the initial distribution.
2. AdaptiveRepart (ParMETIS 3) and Repart\* (ParMETIS 2) - Incremental algorithms with small migration cost.
3. Refine\* - Refines a given partitioning (balance). Can be applied multiple times to reduce the communication cost (cut weight) if desired.

As a rule of thumb, use one of the Part\* methods if you have a poor initial balance and you are willing to spend some time doing migration. One such case is static load balancing; that is, you need to balance only once. Use AdaptiveRepart or the Repart\* methods when you have a reasonably good load balance that you wish to update incrementally. These methods are well suited for dynamic load balancing (for example, adaptive mesh refinement). A reasonable strategy is to call PartKway once to obtain a good initial balance and later update this balance using AdaptiveRepart (Repart\* in ParMetis 2.0).

Zoltan is currently compatible with ParMETIS versions 3.1 and 2.0. There is no guarantee that Zoltan will work correctly if you have a different version of ParMETIS on your computer. (ParMETIS 3.0 will work with Zoltan in most cases, but is not officially supported. ParMETIS 3.1 is highly recommended.



The 2.0 version will soon become obsolete and may not be supported in future Zoltan versions.) The ParMETIS source code can be obtained from the [ParMETIS home page](#). As a courtesy service, a recent, compatible version of the ParMETIS source code is distributed with Zoltan. However, ParMETIS is a completely separate library. If you do not wish to install ParMETIS, it is possible to compile Zoltan without any references to ParMETIS (when you 'make' Zoltan, comment out the `PARMETIS_LIBPATH` variable in the configuration file [Utilities/Config/Config.<platform>](#)).

Note that Zoltan ignores the imbalance tolerance parameter `IMBALANCE_TOL` when ParMETIS 2.0 is used (the default value 1.05 is used), while `IMBALANCE_TOL` works correctly with ParMETIS 3.0. Zoltan supports the multiconstraint feature of ParMETIS 3 through multiple object weights (see [OBJ\\_WEIGHT\\_DIM](#)).

The graph given to Zoltan/ParMETIS must be symmetric. Any self edges (loops) will be ignored. Multiple edges between a pair of vertices is not allowed. All weights must be non-negative. The graph does not have to be connected.

## Method String:

## GRAPH or PARMETIS

### Parameters:

*PARMETIS\_METHOD*

The ParMETIS method to be used; currently nine are available.

*PartKway* - multilevel Kernighan-Lin partitioning

*PartGeom* - space filling curves (coordinate based)

*PartGeomKway* - hybrid method based on PartKway and PartGeom (needs both graph data and coordinates)

*AdaptiveRepart* - adaptive repartitioning (only in ParMETIS 3.0 and higher)

*RepartLDiffusion* - diffusion algorithm (local)

*RepartGDiffusion* - diffusion algorithm (global)

*RepartRemap* - multilevel partitioning with remap seeking to minimize migration cost

*RepartMLRemap* - similar to RepartRemap but with additional multilevel refinement

*RefineKway* - refine the current partitioning (balance)

The method names are case insensitive.

*PARMETIS\_OUTPUT\_LEVEL*

Amount of output the load-balancing algorithm should produce.

0 = no output, 1 = print timing info. Turning on more bits displays more information (for example, 3=1+2, 5=1+4, 7=1+2+4).

*PARMETIS\_COARSE\_ALG*

Coarse algorithm for PartKway. 1 = serial, 2 = parallel. (ParMETIS 2 only)

*PARMETIS\_SEED*

Random seed for ParMETIS.

*PARMETIS\_ITR*

Ratio of interprocessor communication time to redistribution time. A high value will emphasize reducing the edge cut, while a small value will try to keep the change in the new partition (distribution) small. This parameter is only used by AdaptiveRepart. A value of between 100 and 1000 is good for most problems.

*USE\_OBJ\_SIZE*

Use (or not use) the available information about object sizes to estimate migration cost. This parameter is currently only relevant for AdaptiveRepart.

*CHECK\_GRAPH*

Level of error checking for graph input: 0 = no checking, 1 = on-processor checking, 2 = full checking. (CHECK\_GRAPH==2 is very slow and should be used only during debugging).

*SCATTER\_GRAPH*

Scatter graph data by distributing contiguous chunks of objects (vertices) of roughly equal size to each processor before calling the partitioner. 0 = don't scatter; 1 = scatter only if all objects are on a single processor; 2 = scatter if at least one processor owns no objects (recommended to avoid a bug in ParMETIS 2.0); 3 = always scatter.

**Default values:***PARMETIS\_METHOD* = RepartGDiffusion*PARMETIS\_OUTPUT\_LEVEL* = 0*PARMETIS\_COARSE\_ALG* = 2*PARMETIS\_SEED* = 15*PARMETIS\_ITR* = 100*USE\_OBJ\_SIZE* = 1*CHECK\_GRAPH* = 1*SCATTER\_GRAPH* = 1**Required Query Functions:**

For all submethods:

[ZOLTAN\\_NUM\\_OBJ\\_FN](#)[ZOLTAN\\_OBJ\\_LIST\\_FN](#) or[ZOLTAN\\_FIRST\\_OBJ\\_FN/ZOLTAN\\_NEXT\\_OBJ\\_FN](#) pair

Only PartGeom &amp; PartGeomKway:

[ZOLTAN\\_NUM\\_GEOM\\_FN](#)[ZOLTAN\\_GEOM\\_MULTI\\_FN](#) or [ZOLTAN\\_GEOM\\_FN](#)

All but PartGeom:

[ZOLTAN\\_NUM\\_EDGES\\_MULTI\\_FN](#) or[ZOLTAN\\_NUM\\_EDGES\\_FN](#)[ZOLTAN\\_EDGE\\_LIST\\_MULTI\\_FN](#) or[ZOLTAN\\_EDGE\\_LIST\\_FN](#)


---

[\[Table of Contents\]](#) | [\[Next: Jostle\]](#) | [\[Previous: Refinement Tree Partitioning\]](#)

# Graph partitioning: Jostle

[Jostle](#) is a library for graph (mesh) partitioning and load balancing developed at the University of Greenwich, London, UK, by Chris Walshaw [[Jostle](#), [Walshaw](#)]. The parallel version of Jostle is sometimes called pjostle. In the Zoltan context, the name Jostle always refers to the parallel version of the library. The main algorithm used in Jostle is based on multilevel graph partitioning, and a diffusion-type method is available for repartitioning. Hence the Jostle library is very similar to ParMETIS. See the [ParMETIS](#) section for a brief description of graph partitioning as a model for load balancing.

At present, only the most common Jostle options are supported by Zoltan. These are briefly described below. For further details, see the documentation available from the [Jostle home page](#). Other options may be added to Zoltan upon request.

Note that Jostle is not distributed with Zoltan. If you wish to use Jostle within Zoltan, you must first obtain a license for Parallel Jostle and install it on your system. The license is currently free for academic use. Zoltan has been tested only with parallel Jostle version 1.2.\* and may be incompatible with other versions. Zoltan offers only limited support for Jostle and this may be discontinued in the future.

## Method String:

**JOSTLE**

## Parameters:

<i>JOSTLE_OUTPUT_LEVEL</i>	Amount of output Jostle should produce. (integer)
<i>JOSTLE_THRESHOLD</i>	Threshold at which the graph contraction phase is stopped. (integer)
<i>JOSTLE_GATHER_THRESHOLD</i>	Duplicate coarse graph on all processors when there are fewer than this number of nodes. (integer)
<i>JOSTLE_MATCHING</i>	Matching algorithm for graph contraction. (Valid values are "local" and "global".)
<i>JOSTLE_REDUCTION</i>	When reduction is turned off, Jostle performs a diffusion-type algorithm instead of multilevel graph partitioning. (Valid values are "on" and "off".)
<i>JOSTLE_CONNECT</i>	Make a disconnected graph connected before partitioning. (Valid values are "on" and "off".)
<i>CHECK_GRAPH</i>	Level of error checking for graph input: 0 = no checking, 1 = on-processor checking, 2 = full checking. (CHECK_GRAPH==2 is very slow and should be used only during debugging).
<i>SCATTER_GRAPH</i>	Scatter graph data by distributing contiguous chunks of objects (vertices) of roughly equal size to each processor before calling the partitioner. 0 = don't scatter; 1 = scatter only if all objects are on a single processor; 2 = scatter if at least one processor owns no objects; 3 = always scatter.

## Default values:

See the [Jostle](#) documentation. See our [ParMETIS section](#) for the last two parameters.

## Required Query Functions:

[ZOLTAN\\_NUM\\_OBJ\\_FN](#)

[ZOLTAN\\_OBJ\\_LIST\\_FN](#) or

[ZOLTAN\\_FIRST\\_OBJ\\_FN](#)/[ZOLTAN\\_NEXT\\_OBJ\\_FN](#) pair

[ZOLTAN\\_NUM\\_EDGES\\_MULTI\\_FN](#) or

[ZOLTAN\\_NUM\\_EDGES\\_FN](#)

[ZOLTAN\\_EDGE\\_LIST\\_MULTI\\_FN](#) or

[ZOLTAN\\_EDGE\\_LIST\\_FN](#)

---

[\[Table of Contents](#) | [Next: Hypergraph Partitioning](#) | [Previous: ParMETIS\]](#)

# Hypergraph Partitioning

Hypergraph partitioning is a useful partitioning and load balancing method when connectivity data is available. It can be viewed as a more sophisticated alternative to the traditional graph partitioning.

A hypergraph consists of vertices and hyperedges. A hyperedge connects one or more vertices. A graph is a special case of a hypergraph where each edge has size two (two vertices). The hypergraph model is well suited to parallel computing, where vertices correspond to data objects and hyperedges represent the communication requirements. The basic partitioning problem is to partition the vertices into  $k$  approximately equal sets such that the number of cut hyperedges is minimized. Most partitioners (including Zoltan-PHG) allows a more general model where both vertices and hyperedges can be assigned weights. It has been shown that the hypergraph model gives a more accurate representation of communication cost (volume) than the graph model. In particular, for sparse matrix-vector multiplication, the hypergraph model **exactly** represents communication volume. Sparse matrices can be partitioned either along rows or columns; in the row-net model the columns are vertices and each row corresponds to an hyperedge, while in the column-net model the roles of vertices and hyperedges are reversed.

Zoltan contains a native parallel hypergraph partitioner, called PHG (Parallel HyperGraph partitioner). In addition, Zoltan provides access to [PaToH](#), a serial hypergraph partitioner. Note that PaToH is not part of Zoltan and should be obtained separately from the [PaToH web site](#). Zoltan-PHG is a fully parallel multilevel hypergraph partitioner. For further technical description, see [\[Devine et al, 2006\]](#).

Planned future features (currently not supported):

- Fixed vertices
- Repartitioning (to reduce migration cost)
- Multiconstraint partitioning
- 2-dimensional (sparse matrix) partitioning

For applications that already use Zoltan to do graph partitioning, it is fairly easy to upgrade to hypergraph partitioning. See the section [graph vs. hypergraph](#) partitioning.

## Method String:

## HYPERGRAPH

### Parameters:

HYPERGRAPH_PACKAGE	PHG (parallel) or PaToH (serial)
PHG_OUTPUT_LEVEL	Level of verbosity; 0 is silent.
CHECK_HYPERGRAPH	Check that the query functions return valid input data; 0 or 1.
(CHECK_GRAPH)	(This slows performance; intended for debugging.)

*PHG\_COARSENING\_METHOD*

The method to use in the matching/coarsening phase:  
 ipm - inner product matching (a.k.a. heavy connectivity matching); gives best quality.  
 l-ipm - local ipm on each processor. Faster but usually gives poorer quality.  
 a-ipm - alternate between ipm and l-ipm. (A compromise between speed and quality.)  
 none - no coarsening

*PHG\_COARSEPARTITION\_METHOD*

Method to partition the coarsest (smallest) hypergraph; typically done in serial:  
 random - random  
 linear - linear assignment of the vertices (ordered by the user query function)  
 greedy - greedy method based on minimizing cuts  
 auto - automatically select from the above methods (in parallel, the processes will do different methods)

*PHG\_REFINEMENT\_METHOD*

Refinement algorithm:  
 fm - approximate Fiduccia-Mattheyses (FM)  
 no - no refinement

*PHG\_REFINEMENT\_QUALITY*

Knob to control the trade-off between run time and quality. 1 is the recommended (default) setting, >1 gives more refinement (higher quality partitions but longer run time), while <1 gives less refinement (and poorer quality).

*PHG\_RANDOMIZE\_INPUT*

Randomize layout of vertices and hyperedges in internal parallel 2D layout?

Setting this parameter to 1 often reduces Zoltan-PHG execution time.

Operation to be applied to edge weights supplied by different processes for the same hyperedge:

*PHG\_EDGE\_WEIGHT\_OPERATION*

*add* - the hyperedge weight will be the sum of the supplied weights

*max* - the hyperedge weight will be the maximum of the supplied weights

*error* - if the hyperedge weights are not equal, Zoltan will flag an error, otherwise the hyperedge weight will be the value returned by the processes

*PHG\_EDGE\_SIZE\_THRESHOLD*

Ignore hyperedges greater than this fraction times number of vertices.

*ADD\_OBJ\_WEIGHT*

Add implicit vertex (object) weight. This will be in addition to the user-defined weights. (Note: Multi-weight partitioning is not yet supported, so currently only use this option with [obj\\_weight\\_dim=0](#).) Valid values:

"none"

"unit" or "vertices"

"pins" or "nonzeros" or "vertex degree"

**Default values:**

*HYPERGRAPH\_PACKAGE = PHG*  
 CHECK\_HYPERGRAPH=0  
 PHG\_OUTPUT\_LEVEL=0  
*PHG\_COARSENING\_METHOD=ipm*  
*PHG\_COARSEPARTITION\_METHOD=auto*  
 PHG\_REFINEMENT\_METHOD=fm  
*PHG\_REFINEMENT\_QUALITY=1*  
 PHG\_RANDOMIZE\_INPUT=0  
 PHG\_EDGE\_WEIGHT\_OPERATION=max  
 PHG\_EDGE\_SIZE\_THRESHOLD=0.25  
 ADD\_OBJ\_WEIGHT=none

**Required Query Functions:**

[ZOLTAN\\_NUM\\_OBJ\\_FN](#)  
[ZOLTAN\\_OBJ\\_LIST\\_FN](#) or  
[ZOLTAN\\_FIRST\\_OBJ\\_FN/ZOLTAN\\_NEXT\\_OBJ\\_FN](#)  
 pair  
[ZOLTAN\\_HG\\_SIZE\\_CS\\_FN](#)  
[ZOLTAN\\_HG\\_CS\\_FN](#)

**Optional Query Functions:**

[ZOLTAN\\_HG\\_SIZE\\_EDGE\\_WTS\\_FN](#)  
[ZOLTAN\\_HG\\_EDGE\\_WTS\\_FN](#)

It is possible to provide the graph query functions instead of the hypergraph queries, though this is not recommended. If only graph query functions are registered, Zoltan will automatically create a hypergraph from the graph, but this is not equivalent to graph partitioning. In particular, the edge weights will not be accurate.

---

[\[Table of Contents\]](#) | [\[Next: Octree\]](#) | [\[Previous: Jostle\]](#)



# Octree Partitioning (OCTPART)

The Octree Partitioning algorithm is based upon work in load balancing for parallel mesh generation at Rensselaer Polytechnic Institute [[Flaherty, Loy et al.](#)]. It was implemented in Zoltan by Luis Gervasio, Department of Computer Science, Rensselaer Polytechnic Institute, as his summer project in 1998 [[Gervasio](#)]. An octree is a spatial decomposition of the computational domain in which the root of the tree, representing the entire domain, is recursively divided by two in each coordinate direction (producing eight or four "child" octants in 3D or 2D, respectively) until each subregion holds at most an application-specified number of objects. These subregions are represented by the leaves of the octree. The octree data structure is widely used in mesh generation and adaptive mesh refinement [[Baehmann et al.](#), [Shephard and Georges](#)]. The octree resulting from such a spatial decomposition of the domain can be used to partition an application's work [[Edwards](#), [Pilkington and Baden](#), [Warren and Salmon](#)]. To partition an octree, a traversal of the tree is used to define a global ordering on the leaves of the octree. This global ordering is often referred to as a Space-Filling Curve (SFC). The leaves of the octree can be easily assigned to processors in a manner which equally distributes work by assigning slices of the ordered list to processors. Different tree-traversal algorithms produce different global orderings or SFCs, with some SFCs having better connectivity and partition quality properties than others. Currently, Morton Indexing (i.e., Z-curve), Grey Code, and Hilbert SFCs are supported. Morton Indexing and Grey Code SFCs are the simplest (and currently, the fastest) of the SFC algorithms, but they produce lower-quality partitions than the Hilbert SFC.

## Method String:

## OCTPART

### Parameters:

*OCT\_DIM*

Specifies whether the 2D or 3D Octree algorithms should be used. The 3D algorithms can be used for 2D problems, but much memory will be wasted to allow for a non-existent third dimension. Similarly, a 2D algorithm can be used for 3D surface meshes provided that the surface can be projected to the *xy*-plane without overlapping points.

2 = use 2D algorithm; 3 = use 3D algorithm.

*OCT\_METHOD*

The SFC to be used.

0 = Morton Indexing; 1 = Grey Code; 2 = Hilbert.

*OCT\_MINOBJECTS*

The minimum number of objects to allow in a leaf octant of the octree. These objects will be assigned as a group to a processor, so this parameter helps define the granularity of the load-balancing problem. Values greater than or equal to one are allowable.

*OCT\_MAXOBJECTS*

The maximum number of objects to allow in a leaf octant of the octree. These objects will be assigned as a group to a processor, so this parameter helps define the granularity of the load-balancing problem. Values greater than or equal to one are allowable.

*OCT\_OUTPUT\_LEVEL* Amount of output the load-balancing algorithm should produce.  
0 = no statistics; 1 = statistics summary; 2 = debugging information; 3 = data for generating plots.

**Default:**

*OCT\_DIM* = 3  
*OCT\_METHOD* = 2  
*OCT\_MINOBJECTS* = 10  
*OCT\_MAXOBJECTS* = 40  
*OCT\_OUTPUT\_LEVEL* = 0

**Required Query Functions:**

[ZOLTAN\\_NUM\\_OBJ\\_FN](#)  
[ZOLTAN\\_OBJ\\_LIST\\_FN](#) or  
[ZOLTAN\\_FIRST\\_OBJ\\_FN](#)/[ZOLTAN\\_NEXT\\_OBJ\\_FN](#) pair  
[ZOLTAN\\_NUM\\_GEOM\\_FN](#)  
[ZOLTAN\\_GEOM\\_MULTI\\_FN](#) or [ZOLTAN\\_GEOM\\_FN](#)

---

[[Table of Contents](#) | [Next: Ordering](#) | [Previous: Hypergraph Partitioning](#)]

# Ordering Algorithms

The following graph ordering algorithms are currently included in the Zoltan library:

[Nested dissection by METIS/ParMETIS \(NODEND\)](#)

The parenthetical string is the parameter value for [ORDER\\_METHOD](#) parameter; the parameter is set through a call to [Zoltan\\_Set\\_Param](#).

---

## Ordering Parameters

While the overall behavior of Zoltan is controlled by [general Zoltan parameters](#), the behavior of each ordering method is controlled by parameters specific to ordering which are also set by calls to [Zoltan\\_Set\\_Param](#). Many of these parameters are specific to individual ordering algorithms, and are listed in the descriptions of the individual algorithms. However, several have meaning across multiple ordering algorithms. These parameters are described below.

### Parameters:

<i>ORDER_METHOD</i>	The order algorithm used by Zoltan is specified by this parameter. Valid values are <div> "NOEND" (for <a href="#">nodal nested dissection by ParMETIS or METIS</a>),  "METIS" (same as NOEND with ORDER_TYPE = local),  "PARMETIS" (same as NOEND with ORDER_TYPE = global), and  "NONE" (for no load-balancing). </div>
<i>ORDER_TYPE</i>	"LOCAL" or "GLOBAL". If LOCAL is selected, then each processor constructs a local (sub-)graph. All inter-processor edges are simply ignored. The ordering arrays returned, <i>rank</i> and <i>iperm</i> , are local permutation vectors in this case.
<i>ORDER_START_INDEX</i>	The start index for the permutation vectors <i>rank</i> and <i>iperm</i> . Valid values are 0 and 1.
<i>REORDER</i>	If this value is set to TRUE (non-zero), Zoltan assumes that the lists of local and global ids are given as input to <b>Zoltan_Order</b> . Otherwise, the id lists will be populated by <b>Zoltan_Order</b> . The permutation of the ids will be the one produced by calling the query functions.

### Default Values:

*ORDER\_METHOD* = NOEND  
*ORDER\_TYPE* = GLOBAL  
*ORDER\_START\_INDEX* = 0  
*REORDER* = FALSE

---

[[Table of Contents](#) | [Next: Nested dissection by ParMETIS](#) | [Previous: Octtree Algorithm](#)]

# Nested Dissection by METIS/ParMETIS

Nested Dissection (ND) is a popular method to compute fill-reducing orderings for sparse matrices. It can also be used for other ordering purposes. The algorithm recursively finds a separator (bisector) in a graph, orders the nodes in the two subsets first, and nodes in the separator last. In METIS/ParMETIS, the recursion is stopped when the graph is smaller than a certain size, and some version of minimum degree ordering is applied to the remaining graph.

METIS computes a local ordering of the objects on each processor, while ParMETIS performs a global ordering of all the objects. ParMETIS currently (versions 2.0 and 3.0) requires that the number of processors is a power of two.

The generic name for this method is NODEND. If GRAPH\_TYPE=GLOBAL ParMETIS is called, but if it is LOCAL, METIS is called. Alternatively, the user can simply set ORDER\_METHOD to METIS or PARMETIS.

**Order\_Method  
String:**

**NODEND or METIS or PARMETIS**

**Parameters:**

See [ParMETIS](#).

Note that the PARMETIS options are ignored when METIS is called.

**Required Query  
Functions:**

[ZOLTAN\\_NUM\\_OBJ\\_FN](#)

[ZOLTAN\\_OBJ\\_LIST\\_FN](#) or

[ZOLTAN\\_FIRST\\_OBJ\\_FN](#)/[ZOLTAN\\_NEXT\\_OBJ\\_FN](#) pair

[ZOLTAN\\_NUM\\_EDGES\\_MULTI\\_FN](#) or [ZOLTAN\\_NUM\\_EDGES\\_FN](#)

[ZOLTAN\\_EDGE\\_LIST\\_MULTI\\_FN](#) or [ZOLTAN\\_EDGE\\_LIST\\_FN](#)

---

[\[Table of Contents\]](#) | [\[Next: Coloring Algorithms\]](#) | [\[Previous: Ordering Algorithms\]](#)

# Coloring Algorithms

The following coloring algorithms are currently included in the Zoltan library:

[Parallel Coloring](#)

---

## Coloring Parameters

While the overall behavior of Zoltan is controlled by [general Zoltan parameters](#), the behavior of each coloring method is controlled by parameters specific to coloring which are also set by calls to [Zoltan\\_Set\\_Param](#). These parameters are described below.

### Parameters:

<i>DISTANCE</i>	The maximum distance between two objects that should not get the same color is specified by this parameter. Valid values are "1" (for distance-1 coloring) and "2" (for distance-2 coloring).
<i>SUPERSTEP_SIZE</i>	Number of local objects to be colored on each processor before exchanging color information. SUPERSTEP_SIZE should be greater than 0.
<i>COMM_PATTERN</i>	Valid values are "S" (synchronous) and "A" (asynchronous). If synchronous communication is selected, processors are forced to wait for the color information from all other processors to be received before proceeding with coloring of the next SUPERSTEP_SIZE number of local objects. If asynchronous communication is selected, there is no such restriction.
<i>COLOR_ORDER</i>	Valid values are "I" (internal first), "B" (boundary first) and "U" (unordered). If "I" is selected, each processor colors its internal objects before boundary objects. If "B" is selected, each processor colors its boundary objects first. If "U" is selected, there is no such distinction between internal and boundary objects. "U" is not implemented for distance-2 coloring.
<i>COLORING_METHOD</i>	Currently only "F" (first-fit) is implemented. By using "F", the smallest available color that will not cause a conflict is assigned to the object that is being colored.

### Default Values:

*DISTANCE* = 1  
*SUPERSTEP\_SIZE* = 100  
*COMM\_PATTERN* = S  
*COLOR\_ORDER* = I  
*COLORING\_METHOD* = F

---

[Table of Contents](#) | [Next: Parallel Coloring](#) | [Previous: Nested Dissection by ParMETIS](#)

# Parallel Coloring

The parallel coloring algorithm in Zoltan is based on the work of [Boman et al.](#) for distance-1 coloring and [Bozdag et al.](#) for distance-2 coloring. It was implemented in Zoltan by Doruk Bozdag and Umit Catalyurek, Department of Biomedical Informatics, Ohio State University. Distance-1 coloring algorithm is an iterative data parallel algorithm that proceeds in two-phased rounds. In the first phase, processors concurrently color the vertices assigned to them. Adjacent vertices colored in the same parallel step of this phase may result in inconsistencies. In the second phase, processors concurrently check the validity of the colors assigned to their respective vertices and identify a set of vertices that needs to be re-colored in the next round to resolve the detected inconsistencies. The algorithm terminates when every vertex has been colored correctly. To reduce communication frequency, the coloring phase is further decomposed into computation and communication sub-phases. In a communication sub-phase processors exchange recent color information. During a computation sub-phase, a number of vertices determined by the SUPERSTEP\_SIZE parameter, rather than a single vertex, is colored based on currently available color information. With an appropriate choice of a value for SUPERSTEP\_SIZE, the number of ensuing conflicts can be kept low while at the same time preventing the runtime from being dominated by the sending of a large number of small messages. The distance-2 graph coloring problem aims at partitioning the vertex set of a graph into the fewest sets consisting of vertices pairwise at distance greater than two from each other. The algorithm is an extension of the parallel distance-1 coloring algorithm.

## Parameters:

See [Coloring](#)

[Algorithms](#).

## Required Query

## Functions:

[ZOLTAN\\_NUM\\_OBJ\\_FN](#)

[ZOLTAN\\_OBJ\\_LIST\\_FN](#) or

[ZOLTAN\\_FIRST\\_OBJ\\_FN](#)/[ZOLTAN\\_NEXT\\_OBJ\\_FN](#) pair

[ZOLTAN\\_NUM\\_EDGES\\_MULTI\\_FN](#) or [ZOLTAN\\_NUM\\_EDGES\\_FN](#)

[ZOLTAN\\_EDGE\\_LIST\\_MULTI\\_FN](#) or [ZOLTAN\\_EDGE\\_LIST\\_FN](#)

---

[\[Table of Contents\]](#) | [\[Next: Data Services and Utilities\]](#) | [\[Previous: Coloring Algorithms\]](#)



# Data Services and Utilities

Within Zoltan, several utilities are provided to simplify both application development and development of new algorithms in the library. They are separate from the Zoltan library so that applications can use them independently of Zoltan, if desired. They are compiled separately from Zoltan and can be archived in separate libraries. Instructions for [building the utilities](#) and [applications](#) using them are included below; individual library names are listed in the following documentation for each package.

The packages available are listed below.

[Memory Management Utilities](#)

[Unstructured Communication Utilities](#)

[Distributed Directory Utility](#)

---

## Building Utilities

The utilities provided with Zoltan have their own Makefiles and can be built separately from Zoltan. If the user [builds the Zoltan library](#), the utility libraries are built automatically and copied to the appropriate *Zoltan/Obj\_<platform>* directory, where *<platform>* is specified through the [ZOLTAN\\_ARCH environment variable](#). Zoltan and the utilities share the [Utilities/Config/Config.<platform>](#) files specifying compilation paths for various architectures. If, however, a user wishes to use these utilities without using Zoltan, he must build the libraries separately.

The structure and use of Makefiles for the utilities are similar to [Zoltan's makefiles](#); a top-level makefile includes rules for building each utility's library. Object files and the utility libraries are stored in a subdirectory *Obj\_<platform>*, where *<platform>* is a target architecture supported with a [Utilities/Config/Config.<platform>](#) file. The command for compiling a particular utility follows:

```
gmake ZOLTAN_ARCH=<platform> <library_name>
```

where *<library\_name>* is the name of the utility library, and *<platform>* is the target architecture (corresponding to *Utilities/Config/Config.<platform>*). The *<library\_name>* for each utility is included in the following documentation for the utilities.

---

## Building Applications

The utilities are designed so that they can easily be used separately from Zoltan in applications. To enable type-checking of arguments, the function-prototypes file for a utility should be included in all application source code files that directly access the utility. The application must also link with the appropriate utility library (and any other libraries on which the utility depends). Library and function-prototype file names for each utility are listed in the following documentation for the utilities.

---

[\[Table of Contents](#) | [Next: Memory Management Utilities](#) | [Previous: Parallel Coloring\]](#)

# Memory Management Utilities

This package consists of wrappers around the standard C memory allocation and deallocation routines which add error-checking and [debugging capabilities](#). These routines are packaged separately from Zoltan to allow their independent use in other applications. A Fortran90 interface is not yet available. C++ programmers can include the header file "zoltan\_mem.h" and use the C functions. This header file, and in fact all of Zoltan's C language header files, are surrounded by an **extern "C" {}** declaration to prevent name mangling when compiled with a C++ compiler.

<b>Source code location:</b>	<i>Utilities/Memory</i>
<b>Function prototypes file:</b>	<i>Utilities/Memory/zoltan_mem.h or include/zoltan_mem.h</i>
<b>Library name:</b>	libzoltan_mem.a
<b>Other libraries used by this library:</b>	libmpi.a. (See <a href="#">note</a> below.)

**Routines:**

[Zoltan\\_Array\\_Alloc](#): Allocates arrays of dimension  $n$ ,  $n=0,1,\dots,4$

[Zoltan\\_Malloc](#): Wrapper for system malloc.

[Zoltan\\_Calloc](#): Wrapper for system calloc.

[Zoltan\\_Realloc](#): Wrapper for system realloc.

[Zoltan\\_Free](#): Frees memory and sets the pointer to NULL.

[Zoltan\\_Memory\\_Debug](#): Sets the debug level used by the memory utilities; see the [description](#) below.

[Zoltan\\_Memory\\_Stats](#): Prints [memory debugging](#) statistics, such as memory leak information.

[Zoltan\\_Memory\\_Usage](#): Returns user-specified information about memory usage (i.e. maximum memory used, total memory currently allocated).

**Use in Zoltan:**

The memory management utility routines are used extensively in Zoltan and in some individual algorithms. Zoltan developers use these routines directly for most memory management, taking advantage of the error checking and [debugging capabilities](#) of the library.

Rather than call [Zoltan\\_Memory\\_Debug](#) directly, applications using Zoltan can set the [DEBUG\\_MEMORY](#) parameter used by this utility through calls to [Zoltan\\_Set\\_Param](#).

**Note on MPI usage:**

MPI is used only to obtain the processor number (through a call to MPI\_Comm\_rank) for print statements and error messages. If an application does not link with MPI, the memory utilities should be compiled with -DZOLTAN\_NO\_MPI; all output will then appear to be from processor zero, even if it is actually from other processors.

---

```
double *Zoltan_Array_Alloc(char *file, int line, int n, int d1, int d2, ..., int dn, int size);
```

---

The **Zoltan\_Array\_Alloc** routine dynamically allocates an array of dimension  $n$ ,  $n = 0, 1, \dots, 4$  with size ( $d1 \times d2 \times \dots \times dn$ ). It is intended to be used for 2, 3 and 4 dimensional arrays; [Zoltan\\_Malloc](#) should be used for the simpler cases. The memory allocated by **Zoltan\_Array\_Alloc** is contiguous, and can be freed by a single call to [Zoltan\\_Free](#).

### Arguments:

<i>file</i>	A string containing the name of the file calling the function. The <code>__FILE__</code> macro can be passed as this argument. This argument is useful for debugging memory allocation problems.
<i>line</i>	The line number within <i>file</i> of the call to the function. The <code>__LINE__</code> macro can be passed as this argument. This argument is useful for debugging memory allocation problems.
<i>n</i>	The number of dimensions in the array to be allocated. Valid values are 0, 1, 2, 3, or 4.
<i>d1, d2, ..., dn</i>	The size of each dimension to be allocated. One argument is included for each dimension.
<i>size</i>	The size (in bytes) of the data objects to be stored in the array.

### Returned Value:

double *	A pointer to the starting address of the $n$ -dimensional array, or NULL if the allocation fails.
----------	---

### Example:

```
int ** x = (int **) Zoltan_Array_Alloc ( __FILE__ , __LINE__ , 2, 5, 6, sizeof
(int));
```

Allocates a two-dimensional, 5x6-element array of integers.

---

```
double *Zoltan_Malloc(int n, char *file , int line);
```

---

The **Zoltan\_Malloc** function is a wrapper around the standard C malloc routine. It allocates a block of memory of size  $n$  bytes. The principle advantage of using the wrapper is that it allows memory leaks to be tracked via the `DEBUG_MEMORY` variable (set in [Zoltan\\_Memory\\_Debug](#)).

A macro **ZOLTAN\_MALLOC** is defined in *zoltan\_mem.h*. It takes the argument  $n$ , and adds the `__FILE__` and `__LINE__` macros to the argument list of the **Zoltan\_Malloc** call:

```
#define ZOLTAN_MALLOC(n) Zoltan_Malloc((n), __FILE__, __LINE__)
```

Using this macro, the developer gains the file and line debugging information without having to type file and line information in each memory allocation call.

### Arguments:

<i>n</i>	The size (in bytes) of the memory-allocation request.
----------	---

<i>file</i>	A string containing the name of the file calling the function. The <code>__FILE__</code> macro can be passed as this argument. This argument is useful for debugging memory allocation problems.
<i>line</i>	The line number within <i>file</i> of the call to the function. The <code>__LINE__</code> macro can be passed as this argument. This argument is useful for debugging memory allocation problems.

**Returned Value:**

double *	A pointer to the starting address of memory allocated. NULL is returned if $n = 0$ or the routine is unsuccessful.
----------	--

**Example:**

```
struct Zoltan_Struct *b = (struct Zoltan_Struct *)
ZOLTAN_MALLOC(sizeof(struct Zoltan_Struct));
Allocates memory for one Zoltan_Struct data structure.
```

---

```
double *Zoltan_Calloc(int num, int size, char *file, int line);
```

---

The **Zoltan\_Calloc** function is a wrapper around the standard C `calloc` routine. It allocates a block of memory of size *num* \* *size* bytes and initializes the memory to zeros. The principle advantage of using the wrapper is that it allows memory leaks to be tracked via the `DEBUG_MEMORY` variable (set in [Zoltan\\_Set\\_Memory\\_Debug](#)).

A macro **ZOLTAN\_CALLOC** is defined in *zoltan\_mem.h*. It takes the arguments *num* and *size*, and adds the `__FILE__` and `__LINE__` macros to the argument list of the **Zoltan\_Calloc** call:

```
#define ZOLTAN_CALLOC(num, size) Zoltan_Calloc((num), (size), __FILE__,
__LINE__)
```

Using this macro, the developer gains the file and line debugging information without having to type file and line information in each memory allocation call.

**Arguments:**

<i>num</i>	The number of elements of the following <i>size</i> to allocate.
<i>size</i>	The size of each element. Hence, the total allocation is <i>num</i> * <i>size</i> bytes.
<i>file</i>	A string containing the name of the file calling the function. The <code>__FILE__</code> macro can be passed as this argument. This argument is useful for debugging memory allocation problems.
<i>line</i>	The line number within <i>file</i> of the call to the function. The <code>__LINE__</code> macro can be passed as this argument. This argument is useful for debugging memory allocation problems.

**Returned Value:**

double *	A pointer to the starting address of memory allocated. NULL is returned if $n = 0$ or the routine is unsuccessful.
----------	--

**Example:**

```
int *b = (int *) ZOLTAN_CALLOC( 10, sizeof(int));
Allocates memory for 10 integers and initializes the memory to zeros.
```

---

```
double *Zoltan_Realloc(void *ptr, int n, char *file, int line);
```

---

The **Zoltan\_Realloc** function is a "safe" version of realloc. It changes the size of the object pointed to by *ptr* to *n* bytes. The contents of *ptr* are unchanged up to a minimum of the old and new sizes. Error tests ensuring that *n* is a positive number and that space is available to be allocated are performed.

A macro **ZOLTAN\_REALLOC** is defined in *zoltan\_mem.h*. It takes the arguments *ptr* and *n*, and adds the `__FILE__` and `__LINE__` macros to the argument list of the **Zoltan\_Realloc** call:

```
#define ZOLTAN_REALLOC(ptr, n) Zoltan_Realloc((ptr), (n), __FILE__, __LINE__)
```

Using this macro, the developer gains the file and line debugging information without having to type file and line information in each memory allocation call.

### Arguments:

<i>ptr</i>	Pointer to allocated memory to be re-sized.
<i>n</i>	The size (in bytes) of the memory-allocation request.
<i>file</i>	A string containing the name of the file calling the function. The <code>__FILE__</code> macro can be passed as this argument. This argument is useful for debugging memory allocation problems.
<i>line</i>	The line number within <i>file</i> of the call to the function. The <code>__LINE__</code> macro can be passed as this argument. This argument is useful for debugging memory allocation problems.

### Returned Value:

double *	A pointer to the starting address of memory allocated. If the routine is unsuccessful, NULL is returned and <i>*ptr</i> is unchanged.
----------	---

### Example:

```
int n = sizeof(struct Zoltan_Struct);
int *b = (int *) ZOLTAN\_MALLOC (n);
b = (int *) ZOLTAN_REALLOC (b, 2*n);
Reallocates memory for b from length n to length 2*n.
```

---

```
void Zoltan_Free(void **ptr, char *file , int line);
```

---

The **Zoltan\_Free** function calls the system's "free" function for the memory pointed to by *\*ptr*. Note that the argument to this routine has an extra level of indirection when compared to the standard C "free" call. This allows the pointer being freed to be set to NULL, which can help find errors in which a pointer is used after it is deallocated. Error checking is performed to prevent attempts to free NULL pointers. When **Zoltan\_Free** is used with the `DEBUG_MEMORY` options (set in [Zoltan\\_Memory\\_Debug](#)), it can help identify memory leaks.

A macro **ZOLTAN\_FREE** is defined in *zoltan\_mem.h*. It takes the argument *ptr*, and adds the `__FILE__` and `__LINE__` macros to the argument list of the **Zoltan\_Free** call:

```
#define ZOLTAN_FREE(ptr) Zoltan_Free((void **)(ptr), __FILE__, __LINE__)
```

Using this macro, the developer gains the file and line debugging information without having to type file and line information in each memory allocation call.

**Arguments:**

*ptr* Address of a pointer to the memory to be freed. Upon return, *ptr* is set to NULL.

**Example:**

```
ZOLTAN_FREE(& x);
```

Frees memory associated with the variable *x*; upon return, *x* is NULL.

## Debugging Memory Errors

One important reason to use the memory-management utilities' wrappers around the system memory routines is to facilitate debugging of memory problems. Various amounts of information can about memory allocation and deallocation are stored, depending on the debug level set through a call to [Zoltan\\_Memory\\_Debug](#). This information is printed either when an error or warning occurs, or when [Zoltan\\_Memory\\_Stats](#) is called. We have found values of one and two to be very helpful in our development efforts. The routine [Zoltan\\_Memory\\_Usage](#) can be called to return user-specified information about memory utilization to the user's program.

```
void Zoltan_Memory_Debug(int new_level);
```

The **Zoltan\_Memory\_Debug** function sets the level of memory debugging to be used.

**Arguments:**

*new\_level* Integer indicating the amount of debugging to use. Valid options include:

- 0 -- No debugging.
- 1 -- The number of calls to [Zoltan\\_Malloc](#) and [Zoltan\\_Free](#) are tallied, and can be printed by a call to [Zoltan\\_Memory\\_Stats](#).
- 2 -- A list of all calls to [Zoltan\\_Malloc](#) which have not yet been freed is kept. This list is printed by [Zoltan\\_Memory\\_Stats](#) (useful for detecting memory leaks). Any calls to [Zoltan\\_Free](#) with addresses not in this list trigger warning messages. (Note that allocations that occurred prior to setting the debug level to 2 will not be in this list and thus can generate spurious warnings.)
- 3 -- Information about each allocation is printed as it happens.

**Default:**

Memory debug level is 1.

```
void Zoltan_Memory_Stats();
```

The **Zoltan\_Memory\_Stats** function prints information about memory allocation and deallocation. The amount of information printed is determined by the debug level set through a call to [Zoltan\\_Memory\\_Debug](#).



**Arguments:**

None.

---

```
int Zoltan_Memory_Usage(int type);
```

---

The **Zoltan\_Memory\_Usage** function returns information about memory utilization. The memory debug level (set through a call to [Zoltan\\_Set\\_Memory\\_Debug](#)) must be at least 2 for this function to return non-zero values.

**Arguments:***type*

Integer to request type of information required. These integers are defined in *zoltan\_mem.h*. Valid options include:

**ZOLTAN\_MEM\_STAT\_TOTAL** -- The function will return the current total memory allocated via Zoltan's memory allocation routines.

**ZOLTAN\_MEM\_STAT\_MAXIMUM** -- The function will return the maximum total memory allocated via Zoltan's memory allocation routines up to this point.

**Default:**

*type* = *ZOLTAN\_MEM\_STAT\_MAXIMUM*

**Returned Value:**

int

The number in bytes of the specific requested memory statistic.

**Example:**

```
total = Zoltan_Memory_Usage (ZOLTAN_MEM_STAT_TOTAL);
```

---

[[Table of Contents](#) | [Next: Unstructured Communication Utilities](#) | [Previous: Utilities](#)]

# Unstructured Communication Utilities

The unstructured communication package provides a simple interface for doing complicated patterns of point-to-point communication, such as those associated with data remapping. This package consists of a few simple functions which create or modify communication *plans*, perform communication, and destroy communication plans upon completion. The package is descended from software first developed by Steve Plimpton and Bruce Hendrickson, and has proved useful in a variety of different applications. For this reason, it is maintained as a separate library and can be used independently from Zoltan.

In a prototypical usage of the unstructured communication package each processor has some objects to send to other processors, but no processor knows what messages it will receive. A call to [Zoltan\\_Comm\\_Create](#) produces a data structure called a communication *plan* which encapsulates the basic information about the communication operation. The plan does not know anything about the types of objects being transferred, only the number of them. So the same plan can be used repeatedly to transfer different types of data as long as the number of objects in the transfers remains the same. The actual size of objects isn't specified until the call to [Zoltan\\_Comm\\_Do](#) which performs the data transfer.

The plan which is produced by [Zoltan\\_Comm\\_Create](#) assumes that all the objects are of the same size. If this is not true, then a call to [Zoltan\\_Comm\\_Resize](#) can specify the actual size of each object, and the plan is augmented appropriately. [Zoltan\\_Comm\\_Resize](#) can be invoked repeatedly on the same plan to specify varying sizes for different data transfer operations.

Although a friendlier interface may be added in the future, for now the data to be sent must be passed to [Zoltan\\_Comm\\_Do](#) as a packed buffer in which the objects are stored consecutively. This probably requires the application to pull the data out of native data structures and place in into the buffer. The destination of each object is specified by the *proclist* argument to [Zoltan\\_Comm\\_Create](#). Some flexibility is supported by allowing *proclist* to contain negative values, indicating that the corresponding objects are not to be sent. The communication operations allow for any object to be sent to any destination processor. However, if the objects are grouped in such a way that all those being sent to a particular processor are consecutive, the time and memory of an additional copy is avoided.

Function [Zoltan\\_Comm\\_Do\\_Reverse](#) reverses the communication plan to send back messages to the originators.

To allow overlap between communication and processing, POST and WAIT variants of [Zoltan\\_Comm\\_Do](#) and [Zoltan\\_Comm\\_Do\\_Reverse](#) are provided. Communication is initiated by the POST function ([Zoltan\\_Comm\\_Do\\_Post](#) or [Zoltan\\_Comm\\_Do\\_Reverse\\_Post](#)); incoming messages are posted and outgoing messages are sent. Then the user can continue processing. After the processing is complete, the corresponding WAIT function ([Zoltan\\_Comm\\_Do\\_Wait](#) or [Zoltan\\_Comm\\_Do\\_Reverse\\_Wait](#)) is called to wait for all incoming messages to be received. For convenience, these functions use the same calling arguments as [Zoltan\\_Comm\\_Do](#) and [Zoltan\\_Comm\\_Do\\_Reverse](#).

All the functions in the unstructured communication library return integer [error codes](#) identical to those used by Zoltan.

The C++ interface to the unstructured communication utility is found in the `zoltan_comm_cpp.h` header file which defines the **Zoltan\_Comm** class.

A Fortran90 interface is not yet available.

<b>Source code location:</b>	<i>Utilities/Communication</i>
<b>C Function prototypes file:</b>	<i>Utilities/Communication/zoltan_comm.h</i>
<b>C++ class definition:</b>	<i>Utilities/Communication/zoltan_comm_cpp.h</i>
<b>Library name:</b>	libzoltan_comm.a
<b>Other libraries used by this library:</b>	libmpi.a, <a href="#">libzoltan_mem.a</a> .
<b>High Level Routines:</b>	

**[Zoltan\\_Comm\\_Create](#):** computes a communication plan for sending objects to destination processors.

**[Zoltan\\_Comm\\_Do](#):** uses a communication plan to send data objects to destination processors. The POST and WAIT variants are

**[Zoltan\\_Comm\\_Do\\_Post](#)** and

**[Zoltan\\_Comm\\_Do\\_Wait](#).**

**[Zoltan\\_Comm\\_Do\\_Reverse](#):** performs the reverse (opposite) communication of **[Zoltan\\_Comm\\_Do](#)**. The POST and WAIT variants are

**[Zoltan\\_Comm\\_Do\\_Reverse\\_Post](#)** and

**[Zoltan\\_Comm\\_Do\\_Reverse\\_Wait](#).**

**[Zoltan\\_Comm\\_Resize](#):** augments the plan to allow objects to be of variable sizes.

**[Zoltan\\_Comm\\_Copy](#):** create a new communication plan and copy an existing one to it.

**[Zoltan\\_Comm\\_Copy\\_To](#):** copy one existing communication plan to another.

**[Zoltan\\_Comm\\_Destroy](#):** free memory associated with a communication plan.

## Low Level Routines:

**[Zoltan\\_Comm\\_Exchange\\_Sizes](#):** updates the sizes of the messages each processor will receive.

**[Zoltan\\_Comm\\_Invert\\_Map](#):** given a set of messages each processor wants to send, determines the set of messages each processor needs to receive.

**[Zoltan\\_Comm\\_Sort\\_Ints](#):** sorts an array of integer values.

**[Zoltan\\_Comm\\_Info](#):** returns information about a communication plan.

**[Zoltan\\_Comm\\_Invert\\_Plan](#):** given a communication plan, converts the plan into a plan for the reverse communication.

## Use in Zoltan:

The Zoltan library uses the unstructured communication package in its migration tools and in some of the load-balancing algorithms. For example, in [Zoltan\\_Migrate](#), [Zoltan\\_Comm\\_Create](#) is used to develop a communication map for sending objects to be exported to their new destination processors. The sizes of the exported objects are obtained and the communication map is augmented with a call to [Zoltan\\_Comm\\_Resize](#). The data for the objects is packed into a communication buffer and sent to the other processors through a call to [Zoltan\\_Comm\\_Do](#). After the received objects are unpacked, the communication plan is no longer needed, and it is deallocated by a call to [Zoltan\\_Comm\\_Destroy](#). Zoltan developers use the package whenever possible, as improvements made to the package (for example, support for heterogeneous architectures) automatically propagate to the algorithms.

---

**C:**

```
int Zoltan_Comm_Create( struct Zoltan_Comm_Obj **plan, int nsend, int *proclist, MPI_Comm
comm, int tag, int *nreturn);
```

**C++:**

```
Zoltan_Comm( const int & nsend, int *proclist, const MPI_Comm & comm, const int & tag, int
*nreturn);
```

or

```
Zoltan_Comm();
```

```
Zoltan_Comm::Create( const int & nsend, int *proclist, const MPI_Comm & comm, const int & tag,
int *nreturn);
```

---

The **Zoltan\_Comm\_Create** function sets up the communication plan in the unstructured communication package. Its input is a count of objects to be sent to other processors, a list of the processors to which the objects should be sent (repetitions are allowed), and an MPI communicator and tag. It allocates and builds a communication plan that describes to which processors data will be sent and from which processors data will be received. It also computes the amount of data to be sent to and received from each processor. It returns the number of objects to be received by the processor and a pointer to the communication plan it created. The communication plan is then used by calls to [Zoltan\\_Comm\\_Do](#) to perform the actual communication.

## Arguments:

<i>plan</i>	A pointer to the communication plan created by <b>Zoltan_Comm_Create</b> .
<i>nsend</i>	The number of objects to be sent to other processors.
<i>proclist</i>	An array of size <i>nsend</i> of destination processor numbers for each of the objects to be sent.
<i>comm</i>	The MPI communicator for the unstructured communication.
<i>tag</i>	A tag for MPI communication.
<i>nreturn</i>	Upon return, the number of objects to be received by the processor.

## Returned Value:

int	Error code.
-----	-------------

In the C++ interface to the communication utility, the communication plan is represented by a **Zoltan\_Comm** object. It is created when the **Zoltan\_Comm** constructor executes. There are two constructors. The first one listed above uses parameters to initialize the plan. The second constructor does not, but the plan can subsequently be initialized with a call to **Zoltan\_Comm::Create()**.

**C:**

```
int Zoltan_Comm_Do(struct Zoltan_Comm_Obj *plan, int tag, char *send_data, int nbytes, char *recvbuf);
```

```
int Zoltan_Comm_Do_Post(struct Zoltan_Comm_Obj *plan, int tag, char *send_data, int nbytes, char *recvbuf);
```

```
int Zoltan_Comm_Do_Wait(struct Zoltan_Comm_Obj *plan, int tag, char *send_data, int nbytes, char *recvbuf);
```

**C++:**

```
int Zoltan_Comm::Do(const int & tag, char *send_data, const int & nbytes, char *recvbuf);
```

```
int Zoltan_Comm::Do_Post(const int & tag, char *send_data, const int & nbytes, char *recvbuf);
```

```
int Zoltan_Comm::Do_Wait(const int & tag, char *send_data, const int & nbytes, char *recvbuf);
```

The **Zoltan\_Comm\_Do** function performs the communication described in a communication plan built by [Zoltan\\_Comm\\_Create](#). Using the plan, it takes a buffer of object data to be sent and the size (in bytes) of each object's data in that buffer and sends the data to other processors. **Zoltan\_Comm\_Do** also receives object data from other processors and stores it in a receive buffer. The receive buffer must be allocated by the code calling **Zoltan\_Comm\_Do** using the number of received objects returned by [Zoltan\\_Comm\\_Create](#) or [Zoltan\\_Comm\\_Resize](#). If the objects have variable sizes, then [Zoltan\\_Comm\\_Resize](#) must be called before **Zoltan\_Comm\_Do**.

**Arguments:**

<i>plan</i>	A pointer to a communication plan built by <a href="#">Zoltan_Comm_Create</a> .
<i>tag</i>	An MPI message tag.
<i>send_data</i>	A buffer filled with object data to be sent to other processors.
<i>nbytes</i>	The size (in bytes) of the data for one object, or the scale factor if the objects have variable sizes. (See <a href="#">Zoltan_Comm_Resize</a> for more details.)
<i>recvbuf</i>	Upon return, a buffer filled with object data received from other processors.

**Returned Value:**

int	Error code.
-----	-------------

**C:**

```
int Zoltan_Comm_Do_Reverse( struct Zoltan_Comm_Obj *plan, int tag, char *send_data, int nbytes, int *sizes, char *recvbuf);
```

```
int Zoltan_Comm_Do_Reverse_Post( struct Zoltan_Comm_Obj *plan, int tag, char *send_data, int nbytes, int *sizes, char *recvbuf);
```

```
int Zoltan_Comm_Do_Reverse_Wait( struct Zoltan_Comm_Obj *plan, int tag, char *send_data, int nbytes, int *sizes, char *recvbuf);
```

**C++:**

```

int Zoltan_Comm::Do_Reverse( const int & tag, char *send_data, const int & nbytes, int *sizes, char
*recvbuf);
int Zoltan_Comm::Do_Reverse_Post( const int & tag, char *send_data, const int & nbytes, int *sizes,
char *recvbuf);
int Zoltan_Comm::Do_Reverse_Wait( const int & tag, char *send_data, const int & nbytes, int *sizes,
char *recvbuf);

```

---

The **Zoltan\_Comm\_Do\_Reverse** function performs communication based on a communication plan built by [Zoltan\\_Comm\\_Create](#). But unlike [Zoltan\\_Comm\\_Do](#), this routine performs the reverse of the communication pattern. Specifically, all sends in the plan are treated as receives and vice versa. **Zoltan\_Comm\_Do\_Reverse** is particularly well suited to return updated data objects to their originating processors when the objects were initially transferred via [Zoltan\\_Comm\\_Do](#).

### Arguments:

<i>plan</i>	A pointer to a communication plan built by <a href="#">Zoltan_Comm_Create</a> .
<i>tag</i>	An MPI message tag to be used by this routine.
<i>send_data</i>	A buffer filled with object data to be sent to other processors.
<i>nbytes</i>	The size (in bytes) of the data associated with an object, or the scale factor if the objects have variable sizes.
<i>sizes</i>	If not NULL, this input array specifies the size of all the data objects being transferred. This argument is passed directly to <a href="#">Zoltan_Comm_Resize</a> . This array has length equal to the <i>nsend</i> value passed to <b>Zoltan_Comm_Create</b> . But note that for <b>Zoltan_Comm_Do_Reverse</b> this array describes the sizes of the values being received, not sent.
<i>recvbuf</i>	Upon return, a buffer filled with object data received from other processors.

### Returned Value:

int	Error code.
-----	-------------

---

### C:

```

int Zoltan_Comm_Resize( struct Zoltan_Comm_Obj *plan, int *sizes, int tag , int *total_recv_size);

```

### C++:

```

int Zoltan_Comm::Resize( int *sizes, const int & tag , int *total_recv_size);

```

---

If the objects being communicated are of variable sizes, then the plan produced by [Zoltan\\_Comm\\_Create](#) is incomplete. This routine allows the plan to be augmented to allow for variable sizes. **Zoltan\_Comm\_Resize** can be invoked repeatedly on the same plan to specify different object sizes associated with different data transfers.

### Arguments:

<i>plan</i>	A communication plan built by <a href="#">Zoltan_Comm_Create</a> .
-------------	--



<i>sizes</i>	An input array of length equal to the <i>nsend</i> argument in the call to <a href="#">Zoltan_Comm_Create</a> which generated the <i>plan</i> . Each entry in the array is the size of the corresponding object to be sent. If <i>sizes</i> is <i>NULL</i> (on all processors), the objects are considered to be the same size. Note that the true size of a message will be scaled by the <i>nbytes</i> argument to <a href="#">Zoltan_Comm_Do</a> .
<i>tag</i>	A message tag to be used for communication within this routine, based upon the communicator in <i>plan</i> .
<i>total_recv_size</i>	Sum of the sizes of the incoming messages. To get the actual size (in bytes), you need to scale by the <i>nbytes</i> argument to <a href="#">Zoltan_Comm_Do</a> .

**Returned Value:**

int                      Error code.

**C:** struct Zoltan\_Comm\_Obj \***Zoltan\_Comm\_Copy**( struct Zoltan\_Comm\_Obj \**plan*);

**C++:** **Zoltan\_Comm**(const Zoltan\_Comm &*plan*);

**Zoltan\_Comm\_Copy** creates a new Zoltan\_Comm\_Obj structure and copies the existing *plan* to it. The corresponding C++ method is the **Zoltan\_Comm** copy constructor.

**Arguments:**

*plan*                      A pointer to the communication plan to be copied to the new Zoltan\_Comm\_Obj structure.

**Returned Value:**

struct

Zoltan\_Comm\_Obj    the newly created plan, or NULL on error.

\*

**C:** int **Zoltan\_Comm\_Copy\_To**( struct Zoltan\_Comm\_Obj \*\**to*, struct Zoltan\_Comm\_Obj \**from*);

**C++:** Zoltan\_Comm & **operator=**(const Zoltan\_Comm &*plan*);

**Zoltan\_Comm\_Copy\_To** copies one existing communication plan to another. The corresponding C++ method is the **Zoltan\_Comm** copy operator.

**Arguments:**

*to*                        A pointer to a pointer to the communication plan that will be copied to. We destroy the plan first, and set the pointer to the plan to NULL, before proceeding with the copy.

*from*                    A pointer the communication plan that we will make a copy of.

**Returned Value:**

int                      Error code

**C:** int **Zoltan\_Comm\_Destroy**( struct Zoltan\_Comm\_Obj \*\**plan*);

**C++:** ~**Zoltan\_Comm**();



The **Zoltan\_Comm\_Destroy** function frees all memory associated with a communication plan created by [Zoltan\\_Comm\\_Create](#). The C++ **Zoltan\_Comm** object does not have an explicit **Destroy** method. It is deallocated when its destructor is called.

### Arguments:

*plan*                      A pointer to a communication plan built by [Zoltan\\_Comm\\_Create](#). Upon return, *plan* is set to NULL.

### Returned Value:

int                        Error code.

### C:

```
int Zoltan_Comm_Exchange_Sizes( int *sizes_to, int *procs_to, int nsend, int self_msg, int
*sizes_from, int *procs_from, int nrecv, int *total_recv_size, int my_proc, int tag, MPI_Comm comm );
```

### C++:

```
static int Zoltan_Comm::Exchange_Sizes( int *sizes_to, int *procs_to, const int & nsend, const int &
self_msg, int *sizes_from, int *procs_from, const int & nrecv, int *total_recv_size, const int & my_proc,
const int & tag, const MPI_Comm & comm );
```

This routine is used by [Zoltan\\_Comm\\_Resize](#) to update the sizes of the messages each processor is expecting to receive. The processors already know who will send them messages, but if variable sized objects are being communicated, then the sizes of the messages are recomputed and exchanged via this routine.

### Arguments:

*sizes\_to*                      Input array with the size of each message to be sent. Note that the actual number of bytes in the message is the product of this value and the *nbytes* argument to [Zoltan\\_Comm\\_Do](#).

*procs\_to*                      Input array with the destination processor for each of the messages to be sent.

*nsend*                        Input argument with the number of messages to be sent. (Length of the *procs\_to* array.)

*self\_msg*                    Input argument indicating whether a processor has data for itself (=1) or not (=0) within the *procs\_to* and *lengths\_to* arrays.

*sizes\_from*                    Returned array with the size of each message that will be received. Note that the actual number of bytes in the message is the product of this value and the *nbytes* argument to [Zoltan\\_Comm\\_Do](#).

*procs\_from*                    Returned array of processors from which data will be received.

*nrecv*                        Returned value with number of messages to be received. (length of *procs\_from* array.)

*total\_recv\_size*              The total size of all the messages to be received. As above, the actual number of bytes will be scaled by the *nbytes* argument to [Zoltan\\_Comm\\_Do](#).

*my\_proc*                      The processor's ID in the *comm* communicator.

*tag*                          A message tag which can be used by this routine.

*comm*                        MPI Communicator for the processor numbering in the *procs* arrays.

**Returned Value:**

int                      Error code.

**C:**

```
int Zoltan_Comm_Invert_Map( int *lengths_to, int * procs_to, int nsends, int self_msg, int **  
lengths_from, int ** procs_from, int * nrecvs, int my_proc, int nprocs, int out_of_mem, int tag,  
MPI_Comm comm );
```

**C++:**

```
static int Zoltan_Comm::Invert_Map( int *lengths_to, int * procs_to, const int & nsends, const int &  
self_msg, int * & lengths_from, int * & procs_from, int & nrecvs, const int & my_proc, const int &  
nprocs, const int & out_of_mem, const int & tag, const MPI_Comm & comm );
```

The **Zoltan\_Comm\_Invert\_Map** function is a low level communication routine. It is useful when a processor knows to whom it needs to send data, but not from whom it needs to receive data. Each processor provides to this routine a set of lengths and destinations for the messages it wants to send. The routine then returns the set of lengths and origins for the messages a processor will receive. Note that by inverting the interpretation of *to* and *from* in these arguments, the routine can be used to do the opposite: knowing how much data to receive and from which processors, it can compute how much data to send and to which processors.

**Arguments:**

<i>lengths_to</i>	Input array with the number of values in each of the messages to be sent. Note that the actual size of each value is not specified until the <a href="#">Zoltan_Comm_Do</a> routine is invoked.
<i>procs_to</i>	Input array with the destination processor for each of the messages to be sent.
<i>nsends</i>	Input argument with the number of messages to be sent. (Length of the <i>lengths_to</i> and <i>procs_to</i> arrays.)
<i>self_msg</i>	Input argument indicating whether a processor has data for itself (=1) or not (=0) within the <i>procs_to</i> and <i>lengths_to</i> arrays.
<i>lengths_from</i>	Returned array with lengths of messages to be received.
<i>procs_from</i>	Returned array of processors from which data will be received.
<i>nrecvs</i>	Returned value with number of messages to be received (lengths of <i>lengths_from</i> and <i>procs_from</i> arrays).
<i>my_proc</i>	The processor's ID in the <i>comm</i> communicator.
<i>nprocs</i>	Number of processors in the <i>comm</i> communicator.
<i>out_of_mem</i>	Since it has a barrier operation, this routine is a convenient time to tell all the processors that one of them is out of memory. This input argument is 0 if the processor is OK, and 1 if the processor has failed in a malloc call. All the processors will return with a code of <b>COMM_MEMERR</b> if any of them is out of memory.
<i>tag</i>	A message tag which can be used by this routine.
<i>comm</i>	MPI Communicator for the processor numbering in the <i>procs</i> arrays.

**Returned Value:**

int                      Error code.

---

```
int Zoltan_Comm_Sort_Ints( int *vals_sort, int *vals_other, int nvals);
```

---

As its name suggests, the **Zoltan\_Comm\_Sort\_Ints** function sorts a set of integers via the quicksort algorithm. The integers are reordered from lowest to highest, and a second array of integers is reordered in the same fashion. This second array can be used to return the permutation associated with the sort operation. There is no C++ interface to this function. You can use the C function instead.

#### Arguments:

<i>vals_sort</i>	The array of integers to be sorted. This array is permuted into sorted order.
<i>vals_other</i>	Another array of integers which is permuted identically to <i>vals_sort</i> .
<i>nvals</i>	The number of values in the two integer arrays.

#### Returned Value:

int	Error code.
-----	-------------

---

#### C:

```
int Zoltan_Comm_Info( struct Zoltan_Comm_Obj *plan, int *nsends, int *send_procs, int
*send_lengths, int *send_nvals, int *send_max_size, int *send_list, int *nrecvs, int *recv_procs, int
*recv_lengths, int *recv_nvals, int *recv_total_size, int *recv_list, int *self_msg);
```

#### C++:

```
int Zoltan_Comm::Info( int *nsends, int *send_procs, int *send_lengths, int *send_nvals, int
*send_max_size, int *send_list, int *nrecvs, int *recv_procs, int *recv_lengths, int *recv_nvals, int
*recv_total_size, int *recv_list, int *self_msg) const;
```

---

**Zoltan\_Comm\_Info** returns information about a communication plan. All arguments, except the *plan* itself, may be NULL; values are returned only for non-NULL arguments.

#### Arguments:

<i>plan</i>	Communication data structure created by <a href="#">Zoltan_Comm_Create</a> .
<i>nsends</i>	Upon return, the number of processors to which messages are sent; does not include self-messages.
<i>send_procs</i>	Upon return, a list of processors to which messages are sent; self-messages are included.
<i>send_lengths</i>	Upon return, the number of values to be sent to each processor in <i>send_procs</i> .
<i>send_nvals</i>	Upon return, the total number of values to send.
<i>send_max_size</i>	Upon return, the maximum size of a message to be sent; does not include self-messages.
<i>send_list</i>	Upon return, the processor assignment of each value to be sent.
<i>nrecvs</i>	Upon return, the number of processors from which to receive messages; does not include self-messages.
<i>recv_procs</i>	Upon return, a list of processors from which messages are received; includes self-messages.
<i>recv_lengths</i>	Upon return, the number of values to be received from each processor in <i>recv_procs</i> .

<i>recv_nvals</i>	Upon return, the total number of values to receive.
<i>recv_total_size</i>	Upon return, the total size of items to be received.
<i>recv_list</i>	Upon return, the processor assignments of each value to be received.
<i>self_msg</i>	Upon return, the number of self-messages.

**Returned Value:**

int	Error code.
-----	-------------

---

```
C: int Zoltan_Comm_Invert_Plan( struct Zoltan_Comm_Obj **plan );
C++: int Zoltan_Comm::Invert_Plan();
```

---

Given a communication plan, **Zoltan\_Comm\_Invert\_Plan** alters the plan to make it the plan for the reverse communication. Information in the input plan is replaced by information for the reverse-communication plan. All receives in the reverse-communication plan are blocked; thus, using the inverted plan does not produce the same results as [Zoltan\\_Comm\\_Do\\_Reverse](#). If an error occurs within **Zoltan\_Comm\_Invert\_Plan**, the original plan is returned unaltered.

**Arguments:**

<i>plan</i>	Communication data structure created by <a href="#">Zoltan_Comm_Create</a> ; the contents of this plan are irretrievably modified by <b>Zoltan_Comm_Invert_Plan</b> .
-------------	---

**Returned Value:**

int	Error code.
-----	-------------

# Distributed Directory Utility

The owner (i.e. the processor number) of any computational object is subject to change during load balancing. An application may use this directory utility to manage its objects' locations. A distributed directory balances the load (in terms of memory and processing time) and avoids the bottle neck of a centralized directory design.

This distributed directory module may be used alone or in conjunction with Zoltan's load balancing capability and memory and communication services. The user should note that external names (subroutines, etc.) which prefaced by `Zoltan_DD_` are reserved when using this module.

The user initially creates an empty distributed directory using [Zoltan\\_DD\\_Create](#). Then global ID (GID) information is added to the directory using [Zoltan\\_DD\\_Update](#). The directory maintains the GID's basic information: local ID (optional), partition (optional), arbitrary user data (optional), and the current data owner. [Zoltan\\_DD\\_Update](#) is also called after data migration or refinements. [Zoltan\\_DD\\_Find](#) returns the directory information for a list of GIDs. A selected list of GIDs may be removed from the directory by [Zoltan\\_DD\\_Remove](#). When the user has finished using the directory, its memory is returned to the system by [Zoltan\\_DD\\_Destroy](#).

An object is known by its GID. Hashing provides very fast lookup for the information associated with a GID in a two step process. The first hash of the GID yields the processor number owning the directory entry for that GID. The directory entry owner remains constant even if the object (GID) migrates in time. Second, a different hash algorithm of the GID looks up the associated information in directory processor's hash table. The user may optionally register their own (first) hash function to take advantage of their knowledge of their GID naming scheme and the GID's neighboring processors. See the documentation for [Zoltan\\_DD\\_Set\\_Hash\\_Fn](#) for more information. If no user hash function is registered, Zoltan's [Zoltan\\_Hash](#) will be used. This module's design was strongly influenced by the paper "Communication Support for Adaptive Computation" by Pinar and Hendrickson.

Some users number their GIDs by giving the first "n" GIDs to processor 0, the next "n" GIDs to processor 1, and so forth. The function [Zoltan\\_DD\\_Set\\_Neighbor\\_Hash\\_Fn1](#) will provide efficient directory communication when these GIDs stay close to their origin. The function [Zoltan\\_DD\\_Set\\_Neighbor\\_Hash\\_Fn2](#) allows the specification of ranges of GIDs to each processor for more flexibility. The source code for [DD\\_Set\\_Neighbor\\_Hash\\_Fn1](#) and [DD\\_Set\\_Neighbor\\_Hash\\_Fn2](#) provide examples of how a user can create their own "hash" functions taking advantage of their own GID naming convention.

The routine [Zoltan\\_DD\\_Print](#) will print the contents of the directory. The companion routine [Zoltan\\_DD\\_Stats](#) prints out a summary of the hash table size, number of linked lists, and the length of the longest linked list. This may be useful when the user creates their own hash functions.

The C++ interface to this utility is defined in the header file `zoltan_dd_cpp.h` as the class **Zoltan\_DD**. A single **Zoltan\_DD** object represents a distributed directory.

A Fortran90 interface is not yet available.

<b>Source code location:</b>	<i>Utilities/DDirectory</i>
<b>C Function prototypes file:</b>	<i>Utilities/DDirectory/zoltan_dd.h</i>
<b>C++ class definition:</b>	<i>Utilities/DDirectory/zoltan_dd_cpp.h</i>
<b>Library name:</b>	libzoltan_dd.a
<b>Other libraries used by this library:</b>	libmpi.a, libzoltan_mem.a, libzoltan_comm.a

#### **Routines:**

**Zoltan\_DD\_Create**: Allocates memory and initializes the directory.

**Zoltan\_DD\_Copy**: Allocates a new directory structure and copies an existing one to it.

**Zoltan\_DD\_Copy\_To**: Copies one directory structure to another.

**Zoltan\_DD\_Destroy**: Terminate the directory and frees its memory.

**Zoltan\_DD\_Update**: Adds or updates GIDs' directory information.

**Zoltan\_DD\_Find**: Returns GIDs' information (owner, local ID, etc.)

**Zoltan\_DD\_Remove**: Eliminates selected GIDs from the directory.

**Zoltan\_DD\_Stats**: Provides statistics about hash table & linked lists.

**Zoltan\_DD\_Print**: Displays the contents (GIDs, etc) of each directory.

**Zoltan\_DD\_Set\_Hash\_Fn**: Registers a user's optional hash function.

**Zoltan\_DD\_Set\_Neighbor\_Hash\_Fn1**: Hash function with constant number of GIDs per processor.

**Zoltan\_DD\_Set\_Neighbor\_Hash\_Fn2**: Hash function with variable number of GID's per processor.

#### **Data Structures:**

**struct Zoltan\_DD\_Struct**: State & storage used by all DD routines. Users should not modify any internal values in this structure. Users should only pass the address of this structure to the other routines in this package.

#### **C:**

int **Zoltan\_DD\_Create** (struct Zoltan\_DD\_Struct \*\**dd*, MPI\_Comm *comm*, int *num\_gid\_entries*, int *num\_lid\_entries*, int *user\_length*, int *table\_length*, int *debug\_level*);

#### **C++:**

**Zoltan\_DD**( const MPI\_Comm & *comm*, const int & *num\_gid\_entries*, const int & *num\_lid\_entries*, const int & *user\_length*, const int & *table\_length*, const int & *debug\_level*);

or

**Zoltan\_DD**();

**Zoltan\_DD::Create**( const MPI\_Comm & *comm*, const int & *num\_gid\_entries*, const int & *num\_lid\_entries*, const int & *user\_length*, const int & *table\_length*, const int & *debug\_level*);

**Zoltan\_DD\_Create** allocates and initializes memory for the Zoltan\_DD\_Struct structure. It must be called before any other distributed directory routines. MPI must be initialized prior to calling this routine.

The Zoltan\_DD\_Struct must be passed to all other distributed directory routines. The MPI Comm argument designates the processors used for the distributed directory. The MPI Comm argument is



duplicated and stored for later use.

The user can set the debug level argument in the **Zoltan\_DD\_Create** to determine the module's response to multiple updates for any GID within one update cycle. If the argument is set to 0, all multiple updates are ignored (but the last determines the directory information.) If the argument is set to 1, an error is returned if the multiple updates represent different owners for the same GID. If the debug level is 2, an error return and an error message are generated if multiple updates represent different owners for the same GID. If the level is 3, an error return and an error message are generated for a multiple update even if the updates represent the same owner for a GID.

### Arguments:

<i>dd</i>	Structure maintains directory state and hash table.
<i>comm</i>	MPI comm duplicated and stored specifying directory processors.
<i>num_gid_entries</i>	Length of GID.
<i>num_lid_entries</i>	Length of local ID or zero to ignore local IDs.
<i>user_length</i>	Length of user defined data field (optional, may be zero).
<i>table_length</i>	Length of hash table (zero forces default value).
<i>debug_level</i>	Legal values range in [0,3]. Sets response to various error conditions where 3 is the most verbose.

### Returned Value:

int      [Error code](#).

In the C++ interface, the distributed directory is represented by a **Zoltan\_DD** object. It is created when the **Zoltan\_DD** constructor executes. There are two constructors. The first one listed above uses parameters to initialize the distributed directory. The second constructor does not, but it can subsequently be initialized with a call to **Zoltan\_DD::Create()**.

---

```
C: struct Zoltan_DD_Struct *Zoltan_DD_Copy (struct Zoltan_DD_Struct *from);
C++: Zoltan_DD(const Zoltan_DD &dd);
```

---

This routine creates a new distributed directory structure and copies an existing one to it. The corresponding routine in the C++ library is the Zoltan\_DD copy constructor.

### Arguments:

*from*      The existing directory structure which will be copied to the new one.

### Returned Value:

struct  
Zoltan\_DD\_Struct \*      The newly created directory structure.

---

```
C: int Zoltan_DD_Copy_To (struct Zoltan_DD_Struct **to, struct Zoltan_DD_Struct *from);
C++: Zoltan_DD & operator=(const Zoltan_DD &dd);
```

---

This routine copies one distributed directory structure to another. The corresponding method in the C++ library is the Zoltan\_DD class copy operator.



**Arguments:**

<i>to</i>	A pointer to a pointer to the target structure. The structure will be destroyed and the pointer set to NULL before proceeding with the copy.
<i>from</i>	A pointer to the source structure. The contents of this structure will be copied to the target structure.

**Returned Value:**

int	<a href="#">Error code</a> .
-----	------------------------------

**C:** void **Zoltan\_DD\_Destroy** (struct Zoltan\_DD\_Struct \*\**dd*);

**C++:** ~**Zoltan\_DD**();

This routine frees all memory allocated for the distributed directory. No calls to any distributed directory functions using this Zoltan\_DD\_Struct are permitted after calling this routine. MPI is necessary for this routine only to free the previously saved MPI comm.

**Arguments:**

<i>dd</i>	Directory structure to be deallocated.
-----------	--

**Returned Value:**

void	NONE
------	------

There is no explicit **Destroy** method in the C++ **Zoltan\_DD** class. The object is deallocated when it's destructor is called.

**C:**

int **Zoltan\_DD\_Update** (struct Zoltan\_DD\_Struct \**dd*, [ZOLTAN\\_ID\\_PTR](#) *gid*, [ZOLTAN\\_ID\\_PTR](#) *lid*, [ZOLTAN\\_ID\\_PTR](#) *user*, int \**partition*, int *count*);

**C++:**

int **Zoltan\_DD::Update**( [ZOLTAN\\_ID\\_PTR](#) *gid*, [ZOLTAN\\_ID\\_PTR](#) *lid*, [ZOLTAN\\_ID\\_PTR](#) *user*, int \**partition*, const int & *count*);

**Zoltan\_DD\_Update** takes a list of GIDs and corresponding lists of optional local IDs, optional user data, and optional partitions. This routine updates the information for existing directory entries or creates a new entry (filled with given data) if a GID is not found. NULL lists should be passed for optional arguments not desired. This function should be called initially and whenever objects are migrated to keep the distributed directory current.

The user can set the debug level argument in **Zoltan\_DD\_Create** to determine the module's response to multiple updates for any GID within one update cycle.

**Arguments:**

<i>dd</i>	Distributed directory structure state information.
<i>gid</i>	List of GIDs to update (in).
<i>lid</i>	List of corresponding local IDs (optional) (in).
<i>user</i>	List of corresponding user data (optional) (in).

<i>partition</i>	List of corresponding partitions (optional) (in).
<i>count</i>	Number of GIDs in update list.

**Returned Value:**

int	<a href="#">Error code</a> .
-----	------------------------------

**C:**

```
int Zoltan_DD_Find (Zoltan_DD_Directory *dd, ZOLTAN\_ID\_PTR gid, ZOLTAN\_ID\_PTR lid,
ZOLTAN\_ID\_PTR data, int *partition, int count, int *owner);
```

**C++:**

```
int Zoltan_DD::Find( ZOLTAN\_ID\_PTR gid, ZOLTAN\_ID\_PTR lid, ZOLTAN\_ID\_PTR data, int
*partition, const int & count, int *owner) const;
```

Given a list of GIDs, **Zoltan\_DD\_Find** returns corresponding lists of the GIDs' owners, local IDs, partitions, and optional user data. NULL lists must be provided for optional information not being used.

**Arguments:**

<i>dd</i>	Distributed directory structure state information.
<i>gid</i>	List of GIDs whose information is requested.
<i>lid</i>	Corresponding list of local IDs (optional) (out).
<i>data</i>	Corresponding list of user data (optional) (out).
<i>partition</i>	Corresponding list of partitions (optional) (out).
<i>count</i>	Count of GIDs in above list.
<i>owner</i>	Corresponding list of data owners (out).

**Returned Value:**

int	<a href="#">Error code</a> .
-----	------------------------------

**C:**

```
int Zoltan_DD_Remove (struct Zoltan_DD_Struct *dd, ZOLTAN\_ID\_PTR gid, int count);
```

**C++:**

```
int Zoltan_DD::Remove( ZOLTAN\_ID\_PTR gid, const int & count);
```

**Zoltan\_DD\_Remove** takes a list of GIDs and removes all of their information from the distributed directory.

**Arguments:**

<i>dd</i>	Distributed directory structure state information.
<i>gid</i>	List of GIDs to eliminate from the directory.
<i>count</i>	Number of GIDs to be removed.

**Returned Value:**

int	<a href="#">Error code</a> .
-----	------------------------------

**C:**

```
void Zoltan_DD_Set_Hash_Fn (struct Zoltan_DD_Struct *dd, unsigned int (*hash)
(ZOLTAN\_ID\_PTR, int, unsigned int));
```

**C++:**

```
void Zoltan_DD::Set_Hash_Fn( unsigned int (*hash) (ZOLTAN\_ID\_PTR, int, unsigned int));
```

---

Enables the user to register a new hash function for the distributed directory. (If this routine is not called, the default hash function [Zoltan\\_Hash](#) will be used automatically.) This hash function determines which processor maintains the distributed directory entry for a given GID. Inexperienced users do not need this routine.

Experienced users may elect to create their own hash function based on their knowledge of their GID naming scheme. The user's hash function must have calling arguments compatible with [Zoltan\\_Hash](#). Consider that a user has defined a hash function, myhash, as

```
unsigned int myhash(ZOLTAN\_ID\_PTR gid, int length, unsigned int naverage)
{
    return *gid / naverage ; /* GID length assumed to be 1 ; naverage = total_GIDS/nproc */
}
```

Then the call to register this hash function is:

```
Zoltan_DD_Set_Hash (myhash) ;
```

NOTE: This hash function might group the gid's directory information near the gid's owning processor's neighborhood, for an appropriate naming scheme.

**Arguments:**

<i>dd</i>	Distributed directory structure state information.
<i>hash</i>	Name of user's hash function.

**Returned Value:**

void	NONE
------	------

---

**C:**

```
void Zoltan_DD_Stats (struct Zoltan_DD_Struct *dd);
```

**C++:**

```
void Zoltan_DD::Stats() const;
```

---

This routine prints out summary information about the local distributed directory. It includes the hash table length, number of GIDs stored in the local directory, the number of linked lists, and the length of the longest linked list. The debug level (set by an argument to **Zoltan\_DD\_Create** controls this routine's verbosity.

**Arguments:**

<i>dd</i>	Distributed directory structure for state information
-----------	---

**Returned Value:**

void	NONE
------	------

---

```
int Zoltan_DD_Set_Neighbor_Hash_Fn1 (struct Zoltan_DD_Struct *dd, int size);
```

---

This routine associates the first *size* GIDs to proc 0, the next *size* to proc 1, etc. It assumes the GIDs are consecutive numbers. It assumes that GIDs primarily stay near their original owner. The GID length is assumed to be 1. GIDs outside of the range are evenly distributed among the processors via modulo(number of processors). This is a model for the user to develop their own similar routine.

#### Arguments:

<i>dd</i>	Distributed directory structure state information.
<i>size</i>	Number of consecutive GIDs associated with a processor.

#### Returned Value:

int	<a href="#">Error code</a> .
-----	------------------------------

---

```
int Zoltan_DD_Set_Neighbor_Hash_Fn2 (struct Zoltan_DD_Struct *dd, int *proc, int *low, int *high, int n);
```

---

This routine allows the user to specify a beginning and ending GID "numbers" per directory processor. It assumes that GIDs primarily stay near their original owner. It requires that the numbers of high, low, & proc entries are all *n*. It assumes the GID length is 1. It is a model for the user to develop their own similar routine. Users should note the registration of a cleanup routine to free local static memory when the distributed directory is destroyed. GIDs outside the range specified by high and low lists are evenly distributed among the processors via modulo (number of processors).

#### Arguments:

<i>dd</i>	Distributed directory structure state information.
<i>proc</i>	List of processor ids labeling for corresponding high, low value.
<i>low</i>	List of low GID limits corresponding to proc list.
<i>high</i>	List of high GID limits corresponding to proc list.
<i>n</i>	Number of elements in the above lists. Should be number of processors!

#### Returned Value:

int	<a href="#">Error code</a> .
-----	------------------------------

---

```
C: int Zoltan_DD_Print (struct Zoltan_DD_Struct *dd);
```

```
C++: int Zoltan_DD::Print () const;
```

---

This utility displays (to stdout) the entire contents of the distributed directory at one line per GID.

#### Arguments:

<i>dd</i>	Distributed directory structure state information.
-----------	--

#### Returned Value:

int	<a href="#">Error code</a> .
-----	------------------------------

---

## User's Notes

Because Zoltan places no restrictions on the content or length of GIDs, hashing does not guarantee a balanced distribution of objects in the distributed directory. Note also, the worst case behavior of a hash table lookup is very bad (essentially becoming a linear search). Fortunately, the average behavior is very good! The user may specify their own hash function via [Zoltan\\_DD\\_Set\\_Hash\\_Fn](#) to improve performance.

This software module is built on top of the Zoltan Communications functions for efficiency. Improvements to the communications library will automatically benefit the distributed directory.

### FUTURE:

The C99 capability for variable length arrays would significantly simplify many of these following routines. (It eliminates the malloc/free calls for temporary storage. This helps prevent memory leaks.) Other C99 features may also improve code readability. The "inline" capability can potentially improve performance.

The distributed directory should be implemented via threads. However, MPI is not fully thread aware, yet.

---

[\[Table of Contents](#) | [Next: Examples of Zoltan Usage](#) | [Previous: Unstructured Communication Utilities\]](#)

# Examples of Zoltan Usage

Examples for each part of the Zoltan library are provided:

[General use of Zoltan](#)

[Load-balancing calling sequence](#)

[Data migration calling sequences](#)

[Query functions for a simple application](#)

---

[Table of Contents](#) | [Next: General Usage Example](#) | [Previous: Distributed Data Directories](#)

# General Usage Example

An example of general Zoltan usage is included below. This is a C language example. Similar C++ examples may be found in the *examples* directory.

In this example, [Zoltan\\_Initialize](#) is called using the *argc* and *argv* arguments to the main program. Then a pointer to a Zoltan structure is returned by the call to [Zoltan\\_Create](#). In this example, all processors will be used by Zoltan, as **MPI\_COMM\_WORLD** is passed to [Zoltan\\_Create](#) as the communicator.

Several application query functions are then registered with Zoltan through calls to [Zoltan\\_Set\\_Fn](#). Parameters are set through calls to [Zoltan\\_Set\\_Param](#). The application then performs in computations, including making calls to Zoltan functions and utilities.

Before its execution ends, the application frees memory used by Zoltan by calling [Zoltan\\_Destroy](#).

```

/* Initialize the Zoltan library */
struct Zoltan_Struct *zz;
float version;
...
Zoltan\_Initialize(argc, argv, &version);
zz = Zoltan\_Create(MPI_COMM_WORLD);

/* Register query functions. */
Zoltan\_Set\_Fn(zz, ZOLTAN\_NUM\_GEOM\_FN\_TYPE,
               (void (*)()) user_return_dimension, NULL);
Zoltan\_Set\_Fn(zz, ZOLTAN\_GEOM\_FN\_TYPE,
               (void (*)()) user_return_coords, NULL);
Zoltan\_Set\_Fn(zz, ZOLTAN\_NUM\_OBJ\_FN\_TYPE,
               (void (*)()) user_return_num_node, NULL);
Zoltan\_Set\_Fn(zz, ZOLTAN\_OBJ\_LIST\_FN\_TYPE,
               (void (*)()) user_return_owned_nodes, NULL);

/* Set some Zoltan parameters. */
Zoltan\_Set\_Param(zz, "debug_level", "4");

/* Perform application computations, call Zoltan, etc. */
...

/* Free Zoltan data structure before ending application. */
Zoltan\_Destroy (&zz);

```

*Typical calling sequence for general usage of the Zoltan library.*



```

! Initialize the Zoltan library
type(Zoltan_Struct), pointer :: zz
real(Zoltan_FLOAT) version
integer(Zoltan_INT) ierr
...
ierr = Zoltan\_Initialize(version) ! without argc and argv
zz => Zoltan\_Create(MPI_COMM_WORLD)

! Register load-balancing query functions.
! omit data = C NULL
ierr = Zoltan\_Set\_Fn(zz, ZOLTAN\_NUM\_GEOM\_FN\_TYPE,
user_return_dimension)
ierr = Zoltan\_Set\_Fn(zz, ZOLTAN\_GEOM\_FN\_TYPE,
user_return_coords)
ierr = Zoltan\_Set\_Fn(zz, ZOLTAN\_NUM\_OBJ\_FN\_TYPE,
user_return_num_node)
ierr = Zoltan\_Set\_Fn(zz, ZOLTAN\_OBJ\_LIST\_FN\_TYPE,
user_return_owned_nodes)

! Set some Zoltan parameters.
ierr = Zoltan\_Set\_Param(zz, "debug_level", "4")

! Perform application computations, call Zoltan, etc.
...

! Free Zoltan data structure before ending application.
call Zoltan\_Destroy(zz)

```

*Fortran version of general usage example.*

---

[[Table of Contents](#) | [Next: Load-Balancing Example](#) | [Previous: Examples of Library Usage](#)]

# Load-Balancing Example

An example of the typical calling sequence for load balancing using Zoltan in a finite element application is shown in the [figure](#) below. An application first selects a load-balancing algorithm by setting the [LB\\_METHOD](#) parameter with [Zoltan\\_Set\\_Param](#). Next, other parameter values are set by calls to [Zoltan\\_Set\\_Param](#). After some computation, load balancing is invoked by calling [Zoltan\\_LB\\_Partition](#). The results of the load balancing include the number of nodes to be imported and exported to the processor, lists of global and local IDs of the imported and exported nodes, and source and destination processors of the imported and exported nodes. A returned argument of [Zoltan\\_LB\\_Partition](#) is tested to see whether the new decomposition differs from the old one. If the decompositions differ, some sort of data migration is needed to establish the new decomposition; the details of migration are not shown in this [figure](#) but will be addressed in the [migration examples](#). After the data migration is completed, the arrays of information about imported and exported nodes returned by [Zoltan\\_LB\\_Partition](#) are freed by a call to [Zoltan\\_LB\\_Free\\_Part](#).

```
char *lb_method;
int new, num_imp, num_exp, *imp_procs, *exp_procs;
int *imp_to_part, *exp_to_part;
int num_gid_entries, num_lid_entries;
ZOLTAN\_ID\_PTR imp_global_ids, exp_global_ids;
ZOLTAN\_ID\_PTR imp_local_ids, exp_local_ids;

/* Set load-balancing method. */
read_load_balancing_info_from_input_file(&lb_method);
Zoltan\_Set\_Param(zz, "LB_METHOD", lb_method);

/* Reset some load-balancing parameters. */
Zoltan\_Set\_Param(zz, "RCB_Reuse", "TRUE");

/* Perform computations */
...
/* Perform load balancing */
Zoltan\_LB\_Partition(zz, &new, &num_gid_entries, &num_lid_entries,
&num_imp, &imp_global_ids, &imp_local_ids, &imp_procs, &imp_to_part,
&num_exp, &exp_global_ids, &exp_local_ids, &exp_procs, &exp_to_part);
if (new)
    perform_data_migration(...);

/* Free memory allocated for load-balancing results by Zoltan
```

```

library */
Zoltan_LB_Free_Part(&imp_global_ids, &imp_local_ids, &imp_procs,
&imp_to_part);
Zoltan_LB_Free_Part(&exp_global_ids, &exp_local_ids, &exp_procs,
&exp_to_part);
...

```

*Typical calling sequence for performing load balancing with the Zoltan library.*

```

character(len=3) lb_method
logical new
integer(Zoltan_INT) num_imp, num_exp
integer(Zoltan_INT) num_gid_entries, num_lid_entries
integer(Zoltan_INT), pointer :: imp_procs(:), exp_procs(:)
integer(Zoltan_INT), pointer :: imp_global_ids(:),
exp_global_ids(:) ! global IDs
integer(Zoltan_INT), pointer :: imp_local_ids(:),
exp_local_ids(:) ! local IDs
integer(Zoltan_INT) ierr

! Set load-balancing method.
lb_method = "RCB"
ierr = Zoltan_Set_Param(zz, "LB_METHOD", lb_method)

! Reset some load-balancing parameters
ierr = Zoltan_Set_Param(zz, "RCB_Reuse", "TRUE")

! Perform computations
...
! Perform load balancing
ierr =
Zoltan_LB_Partition(zz,new,num_gid_entries,num_lid_entries, &
    num_imp,imp_global_ids,imp_local_ids, &
    imp_procs,imp_to_part, &
    num_exp,exp_global_ids,exp_local_ids, &
    exp_procs,exp_to_part)
if (new) then
    perform_data_migration(...)
endif

! Free memory allocated for load-balancing results by Zoltan
library
ierr = Zoltan_LB_Free_Part(imp_global_ids, imp_local_ids,
imp_procs, imp_to_part);
ierr = Zoltan_LB_Free_Part(exp_global_ids, exp_local_ids,

```

```
exp_procs, exp_to_part);  
...
```

*Fortran version of the load-balancing example.*

---

[Table of Contents](#) | [Next: Migration Examples](#) | [Previous: General Usage Example](#)

# Migration Examples

Data migration using Zoltan's migration tools can be accomplished in two different ways:

[auto-migration](#), or  
[user-guided migration](#).

The choice of migration method depends upon the complexity of the application's data. For some applications, only the objects used in balancing must be migrated; no auxiliary data structures must be moved. Particle simulations are examples of such applications; load balancing is based on the number of particles per processor, and only the particles and their data must be moved to establish the new decomposition. For such applications, Zoltan's auto-migration tools can be used. Other applications, such as finite element methods, perform load balancing on, say, the nodes of the finite element mesh, but nodes that are moved to new processors also need to have their connected elements moved to the new processors, and migrated elements may also need "ghost" nodes (i.e., copies of nodes assigned to other processors) to satisfy their connectivity requirements on the new processor. This complex data migration requires a more user-controlled approach to data migration than the auto-migration capabilities Zoltan can provide.

## Auto-Migration Example

In the [figure](#) below, an example of the load-balancing calling sequence for a particle simulation using Zoltan's auto-migration tools is shown. The application requests auto-migration by turning on the **AUTO\_MIGRATE** option through a call to [Zoltan\\_Set\\_Param](#) and registers functions to pack and unpack a particle's data. During the call to [Zoltan\\_LB\\_Partition](#), Zoltan computes the new decomposition and, using calls to the packing and unpacking query functions, automatically migrates particles to their new processors. The application then frees the arrays returned by [Zoltan\\_LB\\_Partition](#) and can continue computation without having to perform any additional operations for data migration.

```

/* Tell Zoltan to automatically migrate data for the application.
*/
Zoltan_Set_Param(zz, "AUTO_MIGRATE", "TRUE");

/* Register additional functions for packing and unpacking data
*/
/* by migration tools. */
Zoltan_Set_Fn(zz, ZOLTAN_OBJ_SIZE_FN_TYPE,
              (void (*)()) user_return_particle_data_size, NULL);
Zoltan_Set_Fn(zz, ZOLTAN_PACK_OBJ_FN_TYPE,
              (void (*)()) user_pack_particle_data, NULL);
Zoltan_Set_Fn(zz, ZOLTAN_UNPACK_OBJ_FN_TYPE,
              (void (*)()) user_unpack_particle_data, NULL);
...
/* Perform computations */
...
/* Perform load balancing AND automatic data migration! */
Zoltan_LB_Partition(zz, &new, &num_gid_entries, &num_lid_entries,
&num_imp, &imp_global_ids, &imp_local_ids, &imp_procs, &imp_to_part,

```

```

&num_exp,&exp_global_ids,&exp_local_ids,&exp_procs,&exp_to_part);

/* Free memory allocated for load-balancing results by Zoltan */
Zoltan_LB_Free_Part(&imp_global_ids, &imp_local_ids, &imp_procs,
&imp_to_part);
Zoltan_LB_Free_Part(&exp_global_ids, &exp_local_ids, &exp_procs,
&exp_to_part);
...

```

*Typical calling sequence for using the migration tools' auto-migration capability with the dynamic load-balancing tools.*

## User-Guided Migration Example

In the following [figure](#), an example of user-guided migration using Zoltan's migration tools for a finite element application is shown. Several migration steps are needed to completely rebuild the application's data structures for the new decomposition. On each processor, newly imported nodes need copies of elements containing those nodes. Newly imported elements, then, need copies of "ghost" nodes, nodes that are in the element but are assigned to other processors. Each of these entities (nodes, elements, and ghost nodes) can be migrated in separate migration steps using the functions provided in the migration tools. First, the assignment of nodes to processors returned by [Zoltan\\_LB\\_Partition](#) is established. Query functions that pack and unpack nodes are registered and [Zoltan\\_Migrate](#) is called using the nodal decomposition returned from [Zoltan\\_LB\\_Partition](#). [Zoltan\\_Migrate](#) packs the nodes to be exported, sends them to other processors, and unpacks nodes received by a processor. The packing routine *migrate\_node\_pack* includes with each node a list of the element IDs for elements containing that node. The unpacking routine *migrate\_node\_unpack* examines the list of element IDs and builds a list of requests for elements the processor needs but does not already store. At the end of the nodal migration, each processor has a list of element IDs for elements that it needs to support imported nodes but does not already store. Through a call to [Zoltan\\_Invert\\_Lists](#), each processor computes the list of elements it has to send to other processors to satisfy their element requests. Packing and unpacking routines for elements are registered, and [Zoltan\\_Migrate](#) is again used to move element data to new processors. Requests for ghost nodes can be built within the element packing and unpacking routines, and calls to [Zoltan\\_Invert\\_Lists](#) and [Zoltan\\_Migrate](#), with node packing and unpacking, satisfy requests for ghost nodes. In all three phases of migration, the migration tools handle communication; the application is responsible only for packing and unpacking data and for building the appropriate request lists.

```

/* Assume Zoltan returns a decomposition of the */
/* nodes of a finite element mesh. */
Zoltan_LB_Partition(zz,&new,&num_gid_entries,&num_lid_entries,
&num_imp,&imp_global_ids,&imp_local_ids,&imp_procs,&imp_to_part,
&num_exp,&exp_global_ids,&exp_local_ids,&exp_procs,&exp_to_part);

/* Migrate the nodes as directed by the results of Zoltan_LB_Partition. */
/* While unpacking nodes, build list of requests for elements needed */
/* to support the imported nodes.*/
Zoltan_Set_Fn(zz, ZOLTAN_OBJ_SIZE_FN_TYPE,
(void (*)()) migrate_node_size, NULL);
Zoltan_Set_Fn(zz, ZOLTAN_PACK_OBJ_FN_TYPE,
(void (*)()) migrate_pack_node, NULL);
Zoltan_Set_Fn(zz, ZOLTAN_UNPACK_OBJ_FN_TYPE,
(void (*)()) migrate_unpack_node, NULL);

```

```

Zoltan_Migrate(zz,num_import,imp_global_ids,imp_local_ids,imp_procs,imp_to_part,
               num_export,exp_global_ids,exp_local_ids,exp_procs,exp_to_part);

/* Prepare for migration of requested elements. */
Zoltan_Set_Fn(zz, ZOLTAN_PACK_OBJ_FN_TYPE,
              (void (*)()) migrate_pack_element, NULL);
Zoltan_Set_Fn(zz, ZOLTAN_UNPACK_OBJ_FN_TYPE,
              (void (*)()) migrate_unpack_element, NULL);
Zoltan_Set_Fn(zz, ZOLTAN_OBJ_SIZE_FN_TYPE,
              (void (*)()) migrate_element_size, NULL);

/* From the request lists, a processor knows which elements it needs */
/* to support the imported nodes; it must compute which elements to */
/* send to other processors. */
Zoltan_Invert_Lists(zz, Num_Elt_Requests, Elt_Requests_Global_IDs,
                   Elt_Requests_Local_IDs, Elt_Requests_Procs, Elt_Requests_to_Part,
                   &num_tmp_exp, &tmp_exp_global_ids,
                   &tmp_exp_local_ids, &tmp_exp_procs, &tmp_exp_to_part);

/* Processor now knows which elements to send to other processors. */
/* Send the requested elements. While unpacking elements, build */
/* request lists for "ghost" nodes needed by the imported elements. */
Zoltan_Migrate(zz, Num_Elt_Requests, Elt_Requests_Global_IDs,
               Elt_Requests_Local_IDs, Elt_Requests_Procs, Elt_Request_to_Part,
               num_tmp_exp_objs, tmp_exp_global_ids,
               tmp_exp_local_ids, tmp_exp_procs, tmp_exp_to_part);

/* Repeat process for "ghost" nodes. */
...

```

*Typical calling sequence for user-guided use of the migration tools in Zoltan.*

---

[\[Table of Contents\]](#) | [Next: Query-Function Examples](#) | [Previous: Load-Balancing Example](#)



# Query-Function Examples

Examples of query functions provided by a simple application are included below. The general-interface examples include a simple implementation of [ZOLTAN\\_GEOM\\_FN](#) and [ZOLTAN\\_OBJ\\_LIST\\_FN](#) query functions and variants of the simple implementation that exploit local identifiers and data pointers. Migration examples for packing and unpacking objects are also included. Robust error checking is not included in the routines; application developers should include more explicit error checking in their query functions.

## [General Interface Examples](#)

### [Basic example](#)

### [User-defined data pointer](#)

## [Migration Examples](#)

### [Packing and unpacking functions](#)

All the examples use a mesh data structure consisting of nodes in the mesh. these nodes are the objects passed to Zoltan. A node is described by its 3D coordinates and a global ID number that is unique across all processors. The type definitions for the mesh and node data structures used in the examples are included [below](#).

```
/* Node data structure. */
/* A node consists of its 3D coordinates and */
/* an ID number that is unique across all processors. */
struct Node_Type {
    double Coordinates[3];
    int Global_ID_Num;
};

/* Mesh data structure. */
/* Mesh consists of an array of nodes and */
/* the number of nodes owned by the processor. */
struct Mesh_Type {
    struct Node_Type Nodes[MAX_NODES];
    int Number_Owned;
};
```

*Data types for the query-function examples.*

```

! Node data structure.
! A node consists of its 3D coordinates and
! an ID number that is unique across all processors.
type Node_Type
  real(Zoltan_DOUBLE) :: Coordinates(3)
  integer(Zoltan_INT)  :: Global_ID_Num
end type Node_Type

! Mesh data structure.
! Mesh consists of an array of nodes and
! the number of nodes owned by the processor.
type Mesh_Type
  type(Node_Type) :: Nodes(MAX_NODES)
  integer(Zoltan_INT) :: Number_Owned
end type Mesh_Type

```

*Data types for the Fortran query-function examples.*

## General Interface Query Function Examples

In the following examples, [ZOLTAN\\_OBJ\\_LIST\\_FN](#) and [ZOLTAN\\_GEOM\\_FN](#) query functions are implemented for an application using the mesh and node data structures described [above](#). The nodes are the objects passed to Zoltan.

Through a call to [Zoltan\\_Set\\_Fn](#), the function *user\_return\_owned\_nodes* is registered as the [ZOLTAN\\_OBJ\\_LIST\\_FN](#) query function. It returns global and local identifiers for each node owned by a processor.

The function *user\_return\_coords* is registered as a [ZOLTAN\\_GEOM\\_FN](#) query function. Given the global and local identifiers for a node, this function returns the node's coordinates. All the examples exploit the local identifier to quickly locate nodal data. If such an identifier is not available in an application, a search using the global identifier can be performed.

The [Basic Example](#) includes the simplest implementation of the query routines. In the query routines, it uses global application data structures and a local numbering scheme for the local identifiers. The [User-Defined Data Pointer Example](#) uses only local application data structures; this model is useful if the application does not have global data structures or if objects from more than one data structure are to be passed to Zoltan. Differences between the latter example and the Basic Example are shown in **red**.

### Basic Example

In the simplest example, the query functions access the application data through a global data structure (*Mesh*) representing the mesh. In the calls to [Zoltan\\_Set\\_Fn](#), no pointers to application data are registered with the query function (i.e., the *data* pointer is not used). A node's local identifier is an

integer representing the index in the *Mesh.Nodes* array of the node. The local identifier is set to the index's value in *user\_return\_owned\_nodes*. It is used to access the global *Mesh.Nodes* array in *user\_return\_coords*.

```

/* in application's program file */
#include "zoltan.h"

/* Declare a global Mesh data structure. */
struct Mesh_Type Mesh;

main()
{
...
    /* Indicate that local and global IDs are one integer
each. */
    Zoltan_Set_Param(zz, "NUM_GID_ENTRIES", "1");
    Zoltan_Set_Param(zz, "NUM_LID_ENTRIES", "1");

    /* Register query functions. */
    /* Do not register a data pointer with the functions; */
    /* the global Mesh data structure will be used. */
    Zoltan_Set_Fn(zz, ZOLTAN_GEOM_FN_TYPE,
                  (void (*)()) user_return_coords, NULL);
    Zoltan_Set_Fn(zz, ZOLTAN_OBJ_LIST_FN_TYPE,
                  (void (*)()) user_return_owned_nodes, NULL);
...
}

void user_return_owned_nodes(void *data,
    int num_gid_entries, int num_lid_entries,
    ZOLTAN_ID_PTR global_ids, ZOLTAN_ID_PTR local_ids,
    int wgt_dim, float *obj_wgts,
    int *ierr)
{
    int i;
    /* return global node numbers as global_ids. */
    /* return index into Nodes array for local_ids. */
    for (i = 0; i < Mesh.Number_Owned; i++){
        global_ids[i*num_gid_entries] =
Mesh.Nodes[i].Global_ID_Num;
        local_ids[i*num_lid_entries] = i;
    }
    *ierr = ZOLTAN_OK;
}

void user_return_coords(void *data,

```

```

    int num_gid_entries, int num_lid_entries,
    ZOLTAN\_ID\_PTR global_id, ZOLTAN\_ID\_PTR local_id,
    double *geom_vec, int *ierr)
{
    /* use local_id to index into the Nodes array. */
    geom_vec[0] = Mesh.Nodes[local_id[0]].Coordinates[0];
    geom_vec[1] = Mesh.Nodes[local_id[0]].Coordinates[1];
    geom_vec[2] = Mesh.Nodes[local_id[0]].Coordinates[2];
    *ierr = ZOLTAN_OK;
}

```

*Example of general interface query functions (simplest implementation).*

```

! in application's program file

module Global_Mesh_Data
! Declare a global Mesh data structure.
    type(Mesh_Type) :: Mesh
end module

program query_example_1
use zoltan
...
    ! Indicate that local and global IDs are one integer
each.
    ierr = Zoltan\_Set\_Param(zz, "NUM\_GID\_ENTRIES", "1");
    ierr = Zoltan\_Set\_Param(zz, "NUM\_LID\_ENTRIES", "1");

    ! Register query functions.
    ! Do not register a data pointer with the functions;
    ! the global Mesh data structure will be used.
    ierr = Zoltan\_Set\_Fn(zz, ZOLTAN\_GEOM\_FN\_TYPE,
user_return_coords)
    ierr = Zoltan\_Set\_Fn(zz, ZOLTAN\_OBJ\_LIST\_FN\_TYPE,
user_return_owned_nodes)
...
end program

subroutine user_return_owned_nodes(data, &
    num_gid_entries, num_lid_entries, &
    global_ids, local_ids, wgt_dim, obj_wgts, ierr)
use zoltan
use Global_Mesh_Data
integer(Zoltan\_INT) :: data(1) ! dummy declaration, do not use
integer(Zoltan\_INT), intent(in) :: num_gid_entries,

```

```

num_lid_entries
integer(Zoltan_INT), intent(out) :: global_ids(*),
local_ids(*)
integer(Zoltan_INT), intent(in) :: wgt_dim
real(Zoltan_FLOAT), intent(out) :: obj_wgts(*)
integer(Zoltan_INT), intent(out) :: ierr
integer i
    ! return global node numbers as global_ids.
    ! return index into Nodes array for local_ids.
    do i = 1, Mesh%Number_Owned
        global_ids(1+(i-1)*num_gid_entries) = &
            Mesh%Nodes(i)%Global_ID_Num
        local_ids(1+(i-1)*num_lid_entries) = i
    end do
    ierr = ZOLTAN_OK
end subroutine

subroutine user_return_coords(data, num_gid_entries,
num_lid_entries, &
    global_id, local_id, geom_vec, ierr)
use zoltan
use Global_Mesh_Data
integer(Zoltan_INT) :: data(1) ! dummy declaration, do not use
integer(Zoltan_INT), intent(in) :: num_gid_entries,
num_lid_entries
integer(Zoltan_INT), intent(in) :: global_id(*), local_id(*)
real(Zoltan_DOUBLE), intent(out) :: geom_vec(*)
integer(Zoltan_INT), intent(out) :: ierr
    ! use local_id to index into the Nodes array.
    geom_vec(1:3) = Mesh%Nodes(local_id(1))%Coordinates
    ierr = ZOLTAN_OK
end subroutine

```

*Fortran example of general interface query functions (simplest implementation).*

## User-Defined Data Pointer Example

In this example, the address of a local mesh data structure is registered with the query functions for use by those functions. This change eliminates the need for a global mesh data structure in the application. The address of the local data structure is included as an argument in calls to [Zoltan\\_Set\\_Fn](#). This address is then used in *user\_return\_owned\_nodes* and *user\_return\_coords* to provide data for these routines. It is cast to the [Mesh\\_Type](#) data type and accessed with local identifiers as in the [Basic Example](#). Differences between this example and the [Basic Example](#) are shown in red.

This model is useful when the application does not have a global data structure that can be accessed by the query functions. It can also be used for operations on different data structures. For example, if an application had more than one mesh, load balancing could be performed separately on each mesh without

having different query routines for each mesh. Calls to [Zoltan\\_Set\\_Fn](#) would define which mesh should be balanced, and the query routines would access the mesh currently designated by the [Zoltan\\_Set\\_Fn](#) calls.

```

/* in application's program file */
#include "zoltan.h"

main()
{
    /* declare a local mesh data structure. */
    struct Mesh_Type mesh;
    ...
    /* Indicate that local and global IDs are one integer
each. */
    Zoltan_Set_Param(zz, "NUM_GID_ENTRIES", "1");
    Zoltan_Set_Param(zz, "NUM_LID_ENTRIES", "1");

    /* Register query functions. */
    /* Register the address of mesh as the data pointer. */
    Zoltan_Set_Fn(zz, ZOLTAN_GEOM_FN_TYPE,
        (void (*)()) user_return_coords, &mesh);
    Zoltan_Set_Fn(zz, ZOLTAN_OBJ_LIST_FN_TYPE,
        (void (*)()) user_return_owned_nodes, &mesh);
    ...
}

void user_return_owned_nodes(void *data,
    int num_gid_entries, int num_lid_entries,
    ZOLTAN_ID_PTR global_ids, ZOLTAN_ID_PTR local_ids,
    int wgt_dim, float *obj_wgts,
    int *ierr)
{
    int i;
    /* cast data pointer to type Mesh_Type. */
    struct Mesh_Type *ptr = (struct Mesh_Type *) data;

    /* return global node numbers as global_ids. */
    /* return index into Nodes array for local_ids. */
    for (i = 0; i < ptr->Number_Owned; i++) {
        global_ids[i*num_gid_entries] =
ptr->Nodes[i].Global_ID_Num;
        local_ids[i*num_lid_entries] = i;
    }
    *ierr = ZOLTAN_OK;
}

```

```

void user_return_coords(void *data,
    int num_gid_entries, int num_lid_entries,
    ZOLTAN_ID_PTR global_id, ZOLTAN_ID_PTR local_id,
    double *geom_vec, int *ierr)
{
    /* cast data pointer to type Mesh_Type. */
    struct Mesh_Type *ptr = (struct Mesh_Type *) data;

    /* use local_id to address the requested node. */
    geom_vec[0] = ptr->Nodes[local_id[0]].Coordinates[0];
    geom_vec[1] = ptr->Nodes[local_id[0]].Coordinates[1];
    geom_vec[2] = ptr->Nodes[local_id[0]].Coordinates[2];
    *ierr = ZOLTAN_OK;
}

```

*Example of general interface query functions using the application-defined data pointer.*

```

/* included in file zoltan_user_data.f90 */
! User defined data type as wrapper for Mesh
type Zoltan_User_Data_1
    type(Mesh_type), pointer :: ptr
end type Zoltan_User_Data_1

! in application's program file

program query_example_3
use zoltan
! declare a local mesh data structure and a User_Data to point
to it.
type(Mesh_Type), target :: mesh
type(Zoltan_User_Data_1) data
...
! Indicate that local and global IDs are one integer
each.
ierr = Zoltan_Set_Param(zz, "NUM_GID_ENTRIES", "1");
ierr = Zoltan_Set_Param(zz, "NUM_LID_ENTRIES", "1");

! Register query functions.
! Use the User_Data variable to pass the mesh data
data%ptr => mesh
ierr = Zoltan_Set_Fn(zz, ZOLTAN_GEOM_FN_TYPE,
user_return_coords, data)
ierr = Zoltan_Set_Fn(zz, ZOLTAN_OBJ_LIST_FN_TYPE,
    user_return_owned_nodes, data)

```



```

...
end program

subroutine user_return_owned_nodes(data, &
    num_gid_entries, num_lid_entries, &
    global_ids, local_ids, wgt_dim, obj_wgts, ierr)
use zoltan
type(Zoltan_User_Data_1) :: data
integer(Zoltan_INT), intent(in) :: num_gid_entries,
num_lid_entries
integer(Zoltan_INT), intent(out) :: global_ids(*),
local_ids(*)
integer(Zoltan_INT), intent(in) :: wgt_dim
real(Zoltan_FLOAT), intent(out) :: obj_wgts(*)
integer(Zoltan_INT), intent(out) :: ierr
integer i
type(Mesh_Type), pointer :: Mesh

    ! extract the mesh from the User_Data argument
    Mesh => data%ptr

    ! return global node numbers as global_ids.
    ! return index into Nodes array for local_ids.
    do i = 1, Mesh%Number_Owned
        global_ids(1+(i-1)*num_gid_entries) = &
            Mesh%Nodes(i)%Global_ID_Num
        local_ids(1+(i-1)*num_lid_entries) = i
    end do
    ierr = ZOLTAN_OK
end subroutine

subroutine user_return_coords(data, global_id, local_id, &
    geom_vec, ierr)
use zoltan
type(Zoltan_User_Data_1) :: data
integer(Zoltan_INT), intent(in) :: num_gid_entries,
num_lid_entries
integer(Zoltan_INT), intent(in) :: global_id(*), local_id(*)
real(Zoltan_DOUBLE), intent(out) :: geom_vec(*)
integer(Zoltan_INT), intent(out) :: ierr
type(Mesh_Type), pointer :: Mesh

    ! extract the mesh from the User_Data argument
    Mesh => data%ptr

    ! use local_id to index into the Nodes array.
    geom_vec(1:3) = Mesh%Nodes(local_id(1))%Coordinates
    ierr = ZOLTAN_OK

```

```
end subroutine
```

*Fortran example of general interface query functions using the application-defined data pointer.*

## Migration Examples

### Packing and Unpacking Data

Simple migration query functions for the [Basic Example](#) are included [below](#). These functions are used by the migration tools to move nodes among the processors. The functions `user_size_node`, `user_pack_node`, and `user_unpack_node` are registered through calls to [Zoltan\\_Set\\_Fn](#). Query function `user_size_node` returns the size (in bytes) of data representing a single node. Query function `user_pack_node` copies a given node's data into the communication buffer `buf`. Query function `user_unpack_node` copies a data for one node from the communication buffer `buf` into the `Mesh.Nodes` array on its new processor.

These query routines are simple because the application does not dynamically allocate memory for each node. Such dynamic allocation would have to be accounted for in the [ZOLTAN\\_OBJ\\_SIZE\\_FN](#), [ZOLTAN\\_PACK\\_OBJ\\_FN](#), and [ZOLTAN\\_UNPACK\\_OBJ\\_FN](#) routines.

```
main()
{
...
    /* Register migration query functions. */
    /* Do not register a data pointer with the functions; */
    /* the global Mesh data structure will be used. */
    Zoltan_Set_Fn(zz, ZOLTAN_OBJ_SIZE_FN_TYPE,
                  (void (*)()) user_size_node, NULL);
    Zoltan_Set_Fn(zz, ZOLTAN_PACK_OBJ_FN_TYPE,
                  (void (*)()) user_pack_node, NULL);
    Zoltan_Set_Fn(zz, ZOLTAN_UNPACK_OBJ_FN_TYPE,
                  (void (*)()) user_unpack_node, NULL);
...
}

int user_size_node(void *data,
                  int num_gid_entries, int num_lid_entries,
                  ZOLTAN_ID_PTR global_id, ZOLTAN_ID_PTR local_id, int
                  *ierr)
{
    /* Return the size of data associated with one node. */
    /* This case is simple because all nodes have the same size. */
    /*
        *ierr = ZOLTAN_OK;
```

```

        return(sizeof(struct Node_Type));
    }

void user_pack_node(void *data,
    int num_gid_entries, int num_lid_entries,
    ZOLTAN_ID_PTR global_id, ZOLTAN_ID_PTR local_id,
    int dest_proc, int size, char *buf, int *ierr)
{
    /* Copy the specified node's data into buffer buf. */
    struct Node_Type *node_buf = (struct Node_Type *) buf;

    *ierr = ZOLTAN_OK;
    node_buf->Coordinates[0] =
Mesh.Nodes[local_id[0]].Coordinates[0];
    node_buf->Coordinates[1] =
Mesh.Nodes[local_id[0]].Coordinates[1];
    node_buf->Coordinates[2] =
Mesh.Nodes[local_id[0]].Coordinates[2];
    node_buf->Global_ID_Num =
Mesh.Nodes[local_id[0]].Global_ID_Num;
}

void user_unpack_node(void *data, int num_gid_entries,
    ZOLTAN_ID_PTR global_id, int size,
    char *buf, int *ierr)
{
    /* Copy the node data in buf into the Mesh data structure. */
    int i;
    struct Node_Type *node_buf = (struct Node_Type *) buf;

    *ierr = ZOLTAN_OK;
    i = Mesh.Number_Owned;
    Mesh.Number_Owned = Mesh.Number_Owned + 1;
    Mesh.Nodes[i].Coordinates[0] = node_buf->Coordinates[0];
    Mesh.Nodes[i].Coordinates[1] = node_buf->Coordinates[1];
    Mesh.Nodes[i].Coordinates[2] = node_buf->Coordinates[2];
    Mesh.Nodes[i].Global_ID_Num = node_buf->Global_ID_Num;
}

```

*Example of migration query functions for the [Basic Example](#).*

---

[Table of Contents](#) | [Next: FORTRAN Interface](#) | [Previous: Migration Examples](#)

# FORTRAN Interface

The Fortran interface for Zoltan is a Fortran 90 interface designed similar to the Fortran 90 Bindings for OpenGL [[Mitchell](#)]. There is no FORTRAN 77 interface; however, FORTRAN 77 applications can use Zoltan by adding only a few Fortran 90 statements, which are fully explained in the section on [FORTRAN 77](#), provided that vendor-specific extensions are not heavily used in the application. This section describes how to build the Fortran interface into the Zoltan library, how to call Zoltan from Fortran applications, and how to compile Fortran applications that use Zoltan. Note that the capitalization used in this section is for clarity and need not be adhered to in the application code, since Fortran is case insensitive.

[Compiling Zoltan](#)

[Compiling Applications](#)

[FORTRAN API](#)

[FORTRAN 77](#)

[System Specific Remarks](#)

---

[Table of Contents](#) | [Next: FORTRAN--Compiling Zoltan](#) | [Previous: Query-Function Examples](#)

# FORTRAN: Compiling Zoltan

To include the Fortran interface in the Zoltan library, use the **YES\_FORTRAN** parameter in the make statement; for example

```
gmake YES_FORTRAN=1 ZOLTAN\_ARCH=<platform> zoltan
```

Before compiling the library, make sure that the application's [zoltan\\_user\\_data.f90](#) has been placed in the *Zoltan/fort/* directory.

---

[\[Table of Contents](#) | [Next: FORTRAN--Compiling Applications](#) | [Previous: FORTRAN Interface](#) ]

# FORTRAN: Compiling Applications

To compile a Fortran application using the Zoltan library, the module information files must be made available to most compilers during the compilation phase. Module information files are files generated by the compiler to provide module information to program units that **USE** the module. They usually have suffixes like *.mod* or *.M*. The module information files for the modules in the Zoltan library are located in the [Obj\\_<platform>](#) subdirectory. Most Fortran 90 compilers have a compile line flag to specify directories to be searched for module information files, typically "-I"; check the documentation for your compiler. If your compiler does not have such a flag, you will have to copy the module information files to the directory of the application (or use symbolic links).

The Fortran interface is built into the same library file as the rest of Zoltan, which is found during the compiler link phase with *-lzoltan*. Thus an example compilation line would be

```
f90 -I<path to Zoltan>/Obj\_<platform> application.f90 -lzoltan
```

---

[\[Table of Contents\]](#) | [Next: FORTRAN API](#) | [Previous: FORTRAN--Compiling Zoltan](#)

# FORTRAN API

The Fortran interface for each [Zoltan Interface Function](#) and [Application-Registered Query Function](#) is given along with the C interface. This section contains some general information about the design and use of the Fortran interface.

[Names](#)

[Zoltan module](#)

[Numeric types](#)

[Structures](#)

[Global and local IDs](#)

[Query function data](#)

## Names

All procedure, variable, defined constant and structure names are identical to those in the C interface, except that in Fortran they are case insensitive (either upper or lower case letters can be used).

## Zoltan module

MODULE *zoltan* provides access to all entities in Zoltan that are of use to the application, including kind type parameters, named constants, procedures, and derived types. Any program unit (e.g., main program, module, external subroutine) that needs access to an entity from Zoltan must contain the statement

```
USE zoltan
```

near the beginning.

## Numeric types

The correspondence between Fortran and C numeric types is achieved through the use of kind type parameters. In most cases, the default kind for a Fortran type will match the corresponding C type, but this is not guaranteed. To insure portability of the application code, it is highly recommended that the following kind type parameters be used in the declaration of all variables and constants that will be passed to a Zoltan procedure:

C	Fortran
int	INTEGER(KIND=Zoltan_INT)
float	REAL(KIND=Zoltan_FLOAT)
double	REAL(KIND=Zoltan_DOUBLE)

Note that "KIND=" is optional in declaration statements. The kind number for constants can be attached



to the constant, e.g., 1.0\_Zoltan\_DOUBLE.

## Structures

For any struct in the C interface to Zoltan, e.g. [Zoltan\\_Struct](#), there is a corresponding derived type in the Fortran interface. Variables of this type are declared as demonstrated below:

```
TYPE(Zoltan_Struct) :: zz
```

In the Fortran interface, the internal components of the derived type are PRIVATE and not accessible to the application. However, the application simply passes these variables around, and never needs to access the internal components.

## Global and local IDs

While the C implementation uses arrays of unsigned integers to represent [global and local IDs](#), the Fortran interface uses arrays of integers, as unsigned integers are not available in Fortran. Thus, each ID is represented as an array (possibly of size 1) of integers. Applications that use other data types for their IDs can convert between their data types and Zoltan's in the [application-registered query functions](#).

## Query function data

[Zoltan\\_Set\\_Fn](#) allows the application to pass a pointer to data that will subsequently be passed to the query function being registered. From Fortran this is an optional argument, or can be one of several types. In the simplest cases, an intrinsic array containing the data will be sufficient. For these cases, data can be an assumed size array of type INTEGER(Zoltan\_INT), REAL(Zoltan\_FLOAT) or REAL(Zoltan\_DOUBLE). When the argument is omitted in the call to the registration function, a data argument will still be passed to the query function. This should be declared as an assumed size array of type INTEGER(Zoltan\_INT) and never used.

For more complicated situations, the application may need to pass data in a user-defined type. The strong type checking of Fortran does not allow passing an arbitrary type without modifying the Fortran interface for each desired type. So the Fortran interface provides a type to be used for this purpose,

**Zoltan\_User\_Data\_1**. Since different types of data may need to be passed to different query functions, four such types are provided, using the numerals 1, 2, 3 and 4 as the last character in the name of the type. These types are defined by the application in *zoltan\_user\_data.f90*. If not needed, they must be defined, but can be almost empty as in *fort/zoltan\_user\_data.f90*.

The application may use these types in any appropriate way. If desired, it can define these types to contain the application's data and use the type throughout the application. But it is anticipated that in most cases, the desired type already exists in the application, and the **Zoltan\_User\_Data\_x** types will be used as "wrapper types," containing one or more pointers to the existing types. For example,

```
TYPE mesh
    ! an existing data type with whatever defines a mesh
END TYPE mesh
```

```
TYPE Zoltan_User_Data_2
```

```
    TYPE(mesh), POINTER :: ptr
```

```
END TYPE Zoltan_User_Data_2
```

The application would then set the pointer to the data before calling Zoltan\_Set\_Fn:

```
TYPE(mesh) :: meshdata
```

```
TYPE(Zoltan_User_Data_2) :: query_data
```

```
TYPE(Zoltan_Struct) :: zz
```

```
INTEGER(Zoltan_INT), EXTERNAL :: num_obj_func ! not required for module procedures
```

```
query_data%ptr => meshdata
```

```
ierr = Zoltan_Set_Fn(zz,ZOLTAN_NUM_OBJ_FN_TYPE,num_obj_func,query_data)
```

Note that the existing data type must be available when **Zoltan\_User\_Data\_x** is defined. Therefore it must be defined either in *zoltan\_user\_data.f90* or in a module that is compiled before *zoltan\_user\_data.f90* and **USED** by **MODULE** *zoltan\_user\_data*. For an example that uses a wrapper type, see *fdriver/zoltan\_user\_data.f90*.

---

[\[Table of Contents\]](#) | [Next: FORTRAN 77](#) | [Previous: FORTRAN--Compiling Applications](#)

# FORTRAN 77

There is no FORTRAN 77 interface for Zoltan; however, an existing FORTRAN 77 application can be compiled by a Fortran 90 compiler provided it does not use vendor specific extensions (unless the same extensions are supported by the Fortran 90 compiler), and the application can use Zoltan's Fortran 90 interface with a minimal amount of Fortran 90 additions. This section provides details of the Fortran 90 code that must be added.

When building the Zoltan library, use the file `fort/zoltan_user_data.f90` for `zoltan_user_data.f90`. This assumes that DATA in a call to [ZOLTAN\\_SET\\_FN](#) is either omitted (you can omit arguments that are labeled OPTIONAL in the Fortran API) or an array of type INTEGER, REAL or DOUBLE PRECISION (REAL\*4 and REAL\*8 might be acceptable). If a more complicated set of data is required (for example, two arrays), then it should be made available to the query functions through COMMON blocks.

To get access to the interface, each program unit (main program, subroutine or function) that calls a Zoltan routine must begin with the statement

```
USE ZOLTAN
```

and this should be the first statement after the program, subroutine or function statement (before the declarations).

The pointer to the Zoltan structure returned by [ZOLTAN\\_CREATE](#) should be declared as

```
TYPE(ZOLTAN_STRUCT), POINTER :: ZZ
```

(you can use a name other than ZZ if you wish).

To create the structure, use a pointer assignment statement with the call to [ZOLTAN\\_CREATE](#):

```
ZZ => ZOLTAN\_CREATE(COMMUNICATOR)
```

Note that the assignment operator is "`=>`". If ZZ is used in more than one procedure, then put it in a COMMON block. It cannot be passed as an argument unless the procedure interfaces are made "explicit." (Let's not go there.)

The eight import and export arrays passed to [ZOLTAN\\_LB\\_PARTITION](#) (and other procedures) must be pointers. They should be declared as, for example,

```
INTEGER, POINTER :: IMPORT_GLOBAL_IDS(:)
```

Note that the double colon after POINTER is required, and the dimension must be declared as "`(:)`" with a colon. Like ZZ, if they are used in more than one procedure, pass them through a COMMON block, not as an argument.

Except in the unlikely event that the default kinds of intrinsic types do not match the C intrinsic types, you do not have to use the kind type parameters **Zoltan\_INT**, etc. It is also not necessary to include the INTENT attribute in the declarations of the query functions, so they can be simplified to, for example,

```
SUBROUTINE GET_OBJ_LIST(DATA, GLOBAL_IDS, LOCAL_IDS, WGT_DIM,  
OBJ_WGTS, IERR)
```

```
INTEGER DATA(*),GLOBAL_IDS(*),LOCAL_IDS(*),WGT_DIM,IERR  
REAL OBJ_WGTS(*)
```

to be more consistent with a FORTRAN 77 style.

---

[\[Table of Contents](#) | [Next: FORTRAN--System-Specific Remarks](#) | [Previous: FORTRAN API](#)

# FORTRAN: System-Specific Remarks

System-specific details of the FORTRAN interface are included below.

**The mention of specific products, trademarks, or brand names is for purposes of identification only. Such mention is not to be interpreted in any way as an endorsement or certification of such products or brands by the National Institute of Standards and Technology or Sandia National Laboratories. All trademarks mentioned herein belong to their respective owners.**

[MPICH](#)

[Pacific Sierra](#)

[NASoftware](#)

## MPICH

As of version 1.1.2, the MPICH implementation of MPI is not completely "Fortran 90 friendly." Only one problem was encountered during our tests: the reliance on command line arguments. MPICH uses command line arguments during the start-up process, even if the application does not. Command line arguments are not standard in Fortran, so although most compilers offer it as an extension, each compiler has its own method of handling them. The problem arises when one Fortran compiler is specified during the build of MPICH and another Fortran compiler is used for the application. This should not be a problem on systems where there is only one Fortran compiler, or where multiple Fortran compilers are compatible (for example, FORTRAN 77 and Fortran 90 compilers from the same vendor). If your program can get past the call to `MPI_Init`, then you do not have this problem.

To solve this problem, build MPICH in such a way that it does not include the routines for *iargc* and *getarg* (I have been able to do this by using the `-f95nag` flag when configuring MPICH), and then provide your own versions of them when you link the application. Some versions of these routines are provided in *fdriver/farg\_\**.

## Pacific Sierra

Pacific Sierra Research (PSR) Vastf90 is not currently supported due to bugs in the compiler with no known workarounds. It is not known when or if this compiler will be supported.

## NASoftware

N.A.Software FortranPlus is not currently supported due to problems with the query functions. We anticipate that this problem can be overcome, and support will be added soon.

---

[Table of Contents](#) | [Next: C++ Interface](#) | [Previous: FORTRAN 77](#)

# C++ Interface

The C++ interface to the Zoltan library is contained in the header files listed below. Each header file defines one class. Each class represents a Zoltan data structure and the functions that operate on that data structure. The class methods in the header files call functions in the Zoltan C library. So to use the C++ interface from your application, include the appropriate header file and link with the Zoltan C library.

## header file

## class

*include/zoltan\_cpp.h*

**Zoltan**, representing a [load balancing](#) instance

*Utilities/Communication/zoltan\_comm\_cpp.h*

**Zoltan\_Comm**, representing an [unstructured communication](#) instance

*Utilities/DDirectory/zoltan\_dd\_cpp.h*

**Zoltan\_DD**, representing a [distributed directory](#) instance

*Utilities/Timer/zoltan\_timer\_cpp.h*

**Zoltan\_Timer**, representing a timer instance

More detailed information about the interface may be found in the [Zoltan Developer's Guide](#).

Simple examples of the use of the interface may be found in the *examples/CP*P directory. A more complete example is the test driver [zCPPdrive](#). The source code for this test driver is in the *driver* directory.

A note on declaring application registered query functions from a C++ application may be found in the section titled [Application-Registered Query Functions](#).

Two peculiarities of the wrapping of Zoltan with C++ classes are mentioned here:

1. You must call the C language function [Zoltan\\_Initialize](#) before using the C++ interface to the Zoltan library. This function should only be called once. Due to design choices, the C++ interface maintains no global state that is independent of any instantiated objects, so it does not know if the function has been called or not. Therefore, the C++ wrappers do not call [Zoltan\\_Initialize](#) for you.
2. It is preferable to allocate **Zoltan** objects dynamically so you can explicitly delete them before your application exits. (**Zoltan** objects allocated instead on the stack will be deleted automatically at the completion of the scope in which they were created.) The reason is that the **Zoltan** destructor calls `Zoltan_Destroy()`, which makes an MPI call to free the communicator in use by the **Zoltan** object. However the MPI destructor may have been called before the **Zoltan** destructor. In this case you would receive an error while your application is exiting.

This second point is illustrated in the good and bad example below.

```

int main(int argc, char *argv[])
{
    MPI::Init(argc, argv);
    int rank = MPI::COMM_WORLD.Get_rank();
    int size = MPI::COMM_WORLD.Get_size();

    //Initialize the Zoltan library with a C language call
    float version;
    Zoltan_Initialize(argc, argv, &version);

    //Dynamically create Zoltan object.
    Zoltan *zz = new Zoltan(MPI::COMM_WORLD);
    zz->Set_Param("LB_METHOD", "RCB");

    //Several lines of code would follow, working with zz

    //Explicitly delete the Zoltan object
    delete zz;
    MPI::Finalize();
}

```

*Good example, Zoltan object is dynamically allocated and explicitly deleted before exit.*

```

int main(int argc, char *argv[])
{
    Zoltan zz;

    MPI::Init(argc, argv);
    int rank = MPI::COMM_WORLD.Get_rank();
    int size = MPI::COMM_WORLD.Get_size();

    //Initialize the Zoltan library with a C language call
    float version;
    Zoltan_Initialize(argc, argv, &version);

    zz.Set_Param("LB_METHOD", "RCB");

    //Several lines of code would follow, working with zz

    MPI::Finalize();
}

```

*Bad example, the MPI destructor may execute before the Zoltan destructor at process exit.*



[Table of Contents](#) | [Next: Backward Compatibility](#) | [Previous: FORTRAN: System-Specific](#)  
[Remarks Examples](#)

# Backward Compatibility with Previous Versions of Zoltan

As new features have been added to Zoltan, backward compatibility with previous versions of Zoltan has been maintained. Thus, users of previous versions of Zoltan can upgrade to a new version **without changing their application source code**. Modifications to application source code are needed **only** if the applications use new Zoltan functionality.

Enhancements to the Zoltan interface are described below.

[Versions 1.5 and higher](#)

[Versions 1.3 and higher](#)

---

## Backward Compatibility: Versions 1.5 and higher

The ability to generate more partitions than processors was added to Zoltan in version 1.5. Thus, Zoltan's partitioning and migration routines were enhanced to return and use both partition assignments and processor assignments. New interface and query functions were added to support this additional information. All former Zoltan [parameters](#) apply to the new functions as they did to the old; new parameters [NUM\\_GLOBAL\\_PARTITIONS](#) and [NUM\\_LOCAL\\_PARTITIONS](#) apply only to the new functions.

The table below lists the Zoltan function that uses both partition and processor information, along with the analogous function that returns only processor information. Applications requiring only one partition per processor can use either version of the functions.

Function with Partition and Processor info (v1.5 and higher)	Function with only Processor info (v1.3 and higher)
<a href="#">Zoltan_LB_Partition</a>	<a href="#">Zoltan_LB_Balance</a>
<a href="#">Zoltan_LB_Point_PP_Assign</a>	<a href="#">Zoltan_LB_Point_Assign</a>
<a href="#">Zoltan_LB_Box_PP_Assign</a>	<a href="#">Zoltan_LB_Box_Assign</a>
<a href="#">Zoltan_Invert_Lists</a>	<a href="#">Zoltan_Compute_Destinations</a>
<a href="#">Zoltan_Migrate</a>	<a href="#">Zoltan_Help_Migrate</a>
<a href="#">ZOLTAN_PRE_MIGRATE_PP_FN</a>	<a href="#">ZOLTAN_PRE_MIGRATE_FN</a>
<a href="#">ZOLTAN_MID_MIGRATE_PP_FN</a>	<a href="#">ZOLTAN_MID_MIGRATE_FN</a>
<a href="#">ZOLTAN_POST_MIGRATE_PP_FN</a>	<a href="#">ZOLTAN_POST_MIGRATE_FN</a>

To continue using the v1.3 partition functions, no changes to C or Fortran90 applications are needed. Zoltan interfaces from versions earlier than 1.3 are also still supported (see [below](#)), requiring no changes to application programs.

To use the new v1.5 partitioning functions:

- C users must include file *zoltan.h* in their applications and edit their applications to use the appropriate new functions.
- Fortran90 users must put [user-defined data types](#) in *zoltan\_user\_data.f90* and edit their applications to use the appropriate new functions. The new partitioning functions do not work with user-defined data types in *lb\_user\_const.f90*.

## Backward Compatibility: Versions 1.3 and higher

Versions of Zoltan before version 1.3 used a different naming convention for the Zoltan interface and query functions. All functions in Zoltan v.1.3 and above are prefixed with **Zoltan\_**; earlier versions were prefixed with **LB\_**.

**Zoltan versions 1.3 and above maintain backward compatibility with the earlier Zoltan interface.** Thus, applications that used earlier versions of Zoltan can continue using Zoltan **without changing their source code**.

Only two changes are needed to build the application with Zoltan v.1.3 and higher:

- All Zoltan include files are now in directory *Zoltan/include*. Thus, application include paths must point to this directory.  
(Previously, include files were in *Zoltan/lb*.)
- Applications link with Zoltan now by specifying only *-lzoltan*.  
(Previously, applications had to link with *-lzoltan -lzoltan\_comm -lzoltan\_mem*.)

While it is not necessary for application developers to modify their source code to use Zoltan v.1.3 and above, those who want to update their source code should do the following in their application source files:

- Replace Zoltan calls and constants (**LB\_\***) with new names. The new names can be found through the [index below](#).
- C programs: Include file *zoltan.h* instead of *lbi\_const.h*.
- F90 programs: Put [user-defined data types](#) in file *zoltan\_user\_data.f90* instead of *lb\_user\_const.f90*.

### Backward Compatibility Index for Interface and Query Functions

Name in Earlier Zoltan Versions	Name in Zoltan Version 1.3 and higher
LB_BORDER_OBJ_LIST_FN	<a href="#">ZOLTAN_BORDER_OBJ_LIST_FN</a>

LB_Balance	<a href="#">Zoltan LB Balance</a>
LB_Box_Assign	<a href="#">Zoltan LB Box Assign</a>
LB_CHILD_LIST_FN	<a href="#">ZOLTAN_CHILD_LIST_FN</a>
LB_CHILD_WEIGHT_FN	<a href="#">ZOLTAN_CHILD_WEIGHT_FN</a>
LB_COARSE_OBJ_LIST_FN	<a href="#">ZOLTAN_COARSE_OBJ_LIST_FN</a>
LB_Compute_Destinations	<a href="#">Zoltan Compute Destinations</a>
LB_Create	<a href="#">Zoltan Create</a>
LB_Destroy	<a href="#">Zoltan Destroy</a>
LB_EDGE_LIST_FN	<a href="#">ZOLTAN_EDGE_LIST_FN</a>
LB_Eval	<a href="#">Zoltan LB Eval</a>
LB_FIRST_BORDER_OBJ_FN	<a href="#">ZOLTAN_FIRST_BORDER_OBJ_FN</a>
LB_FIRST_COARSE_OBJ_FN	<a href="#">ZOLTAN_FIRST_COARSE_OBJ_FN</a>
LB_FIRST_OBJ_FN	<a href="#">ZOLTAN_FIRST_OBJ_FN</a>
LB_Free_Data	<a href="#">Zoltan LB Free Data</a>
LB_GEOM_FN	<a href="#">ZOLTAN_GEOM_FN</a>
LB_Help_Migrate	<a href="#">Zoltan Help Migrate</a>
LB_Initialize	<a href="#">Zoltan Initialize</a>
LB_MID_MIGRATE_FN	<a href="#">ZOLTAN_MID_MIGRATE_FN</a>
LB_NEXT_BORDER_OBJ_FN	<a href="#">ZOLTAN_NEXT_BORDER_OBJ_FN</a>
LB_NEXT_COARSE_OBJ_FN	<a href="#">ZOLTAN_NEXT_COARSE_OBJ_FN</a>
LB_NEXT_OBJ_FN	<a href="#">ZOLTAN_NEXT_OBJ_FN</a>
LB_NUM_BORDER_OBJ_FN	<a href="#">ZOLTAN_NUM_BORDER_OBJ_FN</a>
LB_NUM_CHILD_FN	<a href="#">ZOLTAN_NUM_CHILD_FN</a>
LB_NUM_COARSE_OBJ_FN	<a href="#">ZOLTAN_NUM_COARSE_OBJ_FN</a>
LB_NUM_EDGES_FN	<a href="#">ZOLTAN_NUM_EDGES_FN</a>

LB_NUM_GEOM_FN	<a href="#">ZOLTAN_NUM_GEOM_FN</a>
LB_NUM_OBJ_FN	<a href="#">ZOLTAN_NUM_OBJ_FN</a>
LB_OBJ_LIST_FN	<a href="#">ZOLTAN_OBJ_LIST_FN</a>
LB_OBJ_SIZE_FN	<a href="#">ZOLTAN_OBJ_SIZE_FN</a>
LB_PACK_OBJ_FN	<a href="#">ZOLTAN_PACK_OBJ_FN</a>
LB_POST_MIGRATE_FN	<a href="#">ZOLTAN_POST_MIGRATE_FN</a>
LB_PRE_MIGRATE_FN	<a href="#">ZOLTAN_PRE_MIGRATE_FN</a>
LB_Point_Assign	<a href="#">Zoltan_LB_Point_Assign</a>
LB_Set_Fn	<a href="#">Zoltan_Set_Fn</a>
LB_Set_<lb_fn_type>_Fn	<a href="#">Zoltan_Set_&lt;zoltan_fn_type&gt;_Fn</a>
LB_Set_Method	<a href="#">Zoltan_Set_Param</a> with parameter <a href="#">LB_METHOD</a>
LB_Set_Param	<a href="#">Zoltan_Set_Param</a>
LB_UNPACK_OBJ_FN	<a href="#">ZOLTAN_UNPACK_OBJ_FN</a>

---

[[Table of Contents](#) | [Next: References](#) | [Previous: C++ Interface](#)]

# References

1. "ALEGRA -- A Framework for Large Strain Rate Physics."  
<http://sherpa.sandia.gov/9231home/alegra/alegra-frame.html>
2. S. Attaway, T. Barragy, K. Brown, D. Gardner, B. Hendrickson, S. Plimpton and C. Vaughan. "Transient Solid Dynamics Simulations on the Sandia/Intel Teraflop Computer." *Proceedings of SC'97*, San Jose, CA, November, 1997. (Finalist for the Gordon Bell Prize.) U. Catalyurek and C. Aykanat, "Hypergraph-partitioning-based decomposition for parallel sparse matrix vector multiplication", *IEEE Trans. Parallel Dist. Systems*, v. 10, no. 7, (1999) pp. 673--693.
3. P. Baehmann, S. Wittchen, M. Shephard, K. Grice, and M. Yerry. "Robust geometrically based automatic two-dimensional mesh generation." *Intl. J. Numer. Meths. Engrg.*, 24 (1987) 1043-1078.
4. E.G. Boman, D. Bozdag, U. Catalyurek, A.H. Gebremedhin and F. Manne. "A Scalable Parallel Graph Coloring Algorithm for Distributed Memory Computers". *Proceedings of Euro-Par'05*, Lisbon, Portugal, August, 2005.
5. D. Bozdag, U. Catalyurek, A.H. Gebremedhin, F. Manne, E.G. Boman and F. Ozguner. "A Parallel Distance-2 Graph Coloring Algorithm for Distributed Memory Computers". *Proceedings of HPCC'05*, Sorrento, Italy, September, 2005.
6. M. Berger and S. Bokhari. "A partitioning strategy for nonuniform problems on multiprocessors." *IEEE Trans. Computers*, C-36 (1987) 570-580.
7. K.D. Devine, E.G. Boman, R. Heaphy, R.H. Bisseling, U.V. Catalyurek. "Parallel Hypergraph Partitioning for Scientific Computing", Proc. of IPDPS'06, Rhodes, Greece, April 2006.
8. K. Devine, G. Hennigan, S. Hutchinson, A. Salinger, J. Shadid, and R. Tuminaro. "High Performance MP Unstructured Finite Element Simulation of Chemically Reacting Flows." *Proceedings of SC'97*, San Jose, CA, November, 1997. (Finalist for the Gordon Bell Prize.)
9. H.C. Edwards. *A parallel infrastructure for scalable adaptive finite element methods and its application to least squares  $C^{\infty}$  collocation*. Ph.D. Dissertation, Univ. of Texas at Austin, May, 1997.
10. J. Flaherty, R. Loy, M. Shephard, B. Szymanski, J. Teresco and L. Ziantz. "Adaptive local refinement with octree load-balancing for the parallel solution of three-dimensional conservation laws." *J. Parallel Distrib. Comput.*, 47 (1998) 139-152.
11. L. Gervasio. "Final Report." Summer project report, Internal Memo, Department 9103, Sandia National Laboratories, August, 1998.
12. B. Hendrickson and K. Devine. "Dynamic load balancing in computational mechanics." *Comp. Meth. Appl. Mech. Engrg.*, v. 184 (#2-4), p. 485-500, 2000.
13. B. Hendrickson and T.G. Kolda. "Partitioning rectangular and structurally nonsymmetric sparse matrices for parallel computation", *SIAM J. on Sci. Comp.*, v. 21, no. 6, 2001, pp. 2048-2072.
14. B. Hendrickson and R. Leland. "The Chaco user's guide, version 2.0." Tech. Rep. SAND 94-2692, Sandia National Laboratories, Albuquerque, NM, October, 1994.  
<http://www.cs.sandia.gov/CRF/chac.html>

15. G. Karypis and V. Kumar. "ParMETIS: Parallel graph partitioning and sparse matrix ordering library." Tech. Rep. 97-060, Department of Computer Science, Univ. of Minnesota, 1997.  
<http://www-users.cs.umn.edu/~karypis/metis/parmetis/>
16. R. Loy. *Adaptive local refinement with octree load-balancing for the parallel solution of three-dimensional conservation laws*. Ph. D. Dissertation, Dept. of Computer Science, Rensselaer Polytechnic Institute, May 1998.
17. S. Mitchell and S. Vavasis. "Quality mesh generation in three dimensions." *Proc. 8th ACM Symposium on Computational Geometry*, ACM (1992) 212-221.
18. W. F. Mitchell. "A Fortran 90 Interface for OpenGL: Revised January 1998" NISTIR 6134 (1998).  
<http://math.nist.gov/~mitchell/papers/nistir6134.ps.gz>
19. W. F. Mitchell. "The K-way Refinement Tree Partitioning Method for Adaptive Grids."  
<http://math.nist.gov/~mitchell/papers/parcomp.ps.gz>
20. "MPSalsa: Massively Parallel Numerical Methods for Advanced Simulation of Chemically Reacting Flows." <http://www.cs.sandia.gov/CRF/MPSalsa/>
21. A. Patra and J. T. Oden. "Problem decomposition for adaptive hp-finite element methods." *J. Computing Systems in Engrg.*, 6 (1995).
22. J. Pilkington and S. Baden. "Partitioning with space-filling curves." Tech. Rep. CS94-349, Dept. of Computer Science and Engineering, Univ. of California, San Diego, CA, 1994.
23. M. Shephard and M. Georges. "Automatic three-dimensional mesh generation by the finite octree technique." *Intl. J. Numer. Meths. Engrg.*, 32 (1991) 709-749.
24. V. E. Taylor and B. Nour-Omid. "A Study of the Factorization Fill-in for a Parallel Implementation of the Finite Element Method." *Intl. J. Numer. Meths. Engrg.*, 37 (1994) 3809-3823.
25. C. Walshaw. "JOSTLE mesh partitioning software", <http://www.gre.ac.uk/jostle/>
26. C. Walshaw, M. Cross, and M. Everett. "Parallel Dynamic Graph Partitioning for Adaptive Unstructured Meshes", *J. Par. Dist. Comp.*, 47(2) 102-108, 1997.
27. M. Warren and J. Salmon. "A parallel hashed octree n-body algorithm." *Proc. Supercomputing '93*, Portland, OR, November 1993.
28. R. D. Williams. "Performance of dynamic load balancing algorithms for unstructured mesh calculations. *Concurrency, Practice, and Experience*, 3(5), 457-481, 1991.

---

[\[Table of Contents](#) | [Next: Index of Interface and Query Functions](#) | [Previous: Backward Compatibility\]](#)



# Index of Interface and Query Functions

[ZOLTAN\\_BORDER\\_OBJ\\_LIST\\_FN](#)

[ZOLTAN\\_CHILD\\_LIST\\_FN](#)

[ZOLTAN\\_CHILD\\_WEIGHT\\_FN](#)

[ZOLTAN\\_COARSE\\_OBJ\\_LIST\\_FN](#)

[Zoltan\\_Color](#)

[Zoltan\\_Compute\\_Destinations](#)

[Zoltan\\_Create](#)

[Zoltan\\_Destroy](#)

[ZOLTAN\\_EDGE\\_LIST\\_FN](#)

[ZOLTAN\\_EDGE\\_LIST\\_MULTI\\_FN](#)

[ZOLTAN\\_FIRST\\_BORDER\\_OBJ\\_FN](#)

[ZOLTAN\\_FIRST\\_COARSE\\_OBJ\\_FN](#)

[ZOLTAN\\_FIRST\\_OBJ\\_FN](#)

[ZOLTAN\\_GEOM\\_FN](#)

[ZOLTAN\\_GEOM\\_MULTI\\_FN](#)

[Zoltan\\_Help\\_Migrate](#)

[Zoltan\\_Initialize](#)

[Zoltan\\_Invert\\_Lists](#)

[Zoltan\\_LB\\_Balance](#)

[Zoltan\\_LB\\_Box\\_Assign](#)

[Zoltan\\_LB\\_Box\\_PP\\_Assign](#)

[Zoltan\\_LB\\_Eval](#)

[Zoltan\\_LB\\_Free\\_Data](#)

[Zoltan\\_LB\\_Partition](#)

[Zoltan\\_LB\\_Point\\_Assign](#)

[Zoltan\\_LB\\_Point\\_PP\\_Assign](#)

[Zoltan\\_LB\\_Set\\_Part\\_Sizes](#)

[ZOLTAN\\_MID\\_MIGRATE\\_FN](#)

[ZOLTAN\\_MID\\_MIGRATE\\_PP\\_FN](#)

[Zoltan\\_Migrate](#)

[ZOLTAN\\_NEXT\\_BORDER\\_OBJ\\_FN](#)

[ZOLTAN\\_NEXT\\_COARSE\\_OBJ\\_FN](#)

[ZOLTAN\\_NEXT\\_OBJ\\_FN](#)

[ZOLTAN\\_NUM\\_BORDER\\_OBJ\\_FN](#)

[ZOLTAN\\_NUM\\_CHILD\\_FN](#)

[ZOLTAN\\_NUM\\_COARSE\\_OBJ\\_FN](#)  
[ZOLTAN\\_NUM\\_EDGES\\_FN](#)  
[ZOLTAN\\_NUM\\_EDGES\\_MULTI\\_FN](#)  
[ZOLTAN\\_NUM\\_GEOM\\_FN](#)  
[ZOLTAN\\_NUM\\_OBJ\\_FN](#)  
[ZOLTAN\\_OBJ\\_LIST\\_FN](#)  
[ZOLTAN\\_OBJ\\_SIZE\\_FN](#)  
[Zoltan\\_Order](#)  
[ZOLTAN\\_PACK\\_OBJ\\_FN](#)  
[ZOLTAN\\_PARTITION\\_FN](#)  
[ZOLTAN\\_PARTITION\\_MULTI\\_FN](#)  
[ZOLTAN\\_POST\\_MIGRATE\\_FN](#)  
[ZOLTAN\\_POST\\_MIGRATE\\_PP\\_FN](#)  
[ZOLTAN\\_PRE\\_MIGRATE\\_FN](#)  
[ZOLTAN\\_PRE\\_MIGRATE\\_PP\\_FN](#)  
[Zoltan\\_RCB\\_Box](#)  
[Zoltan\\_Set\\_Fn](#)  
[Zoltan\\_Set <zoltan\\_fn\\_type>\\_Fn](#)  
[Zoltan\\_Set\\_Param](#)  
[ZOLTAN\\_UNPACK\\_OBJ\\_FN](#)

---

[\[Table of Contents](#) | [Previous: References](#) | [Zoltan Home Page\]](#)